arXiv:1411.4584v1 [cs.CC] 17 Nov 2014

Pseudorandomness for concentration bounds and signed majorities

Parikshit Gopalan
Microsoft

Daniel M. Kane
University of California, San Diego

Raghu Meka

Abstract

The problem of constructing pseudorandom generators that fool halfspaces has been studied intensively in recent times. For fooling halfspaces over $\{\pm 1\}^n$ with polynomially small error, the best construction known requires seed-length $O(\log^2(n))$ [MZ13]. Getting the seed-length down to $O(\log(n))$ is a natural challenge in its own right, which needs to be overcome in order to derandomize RL. In this work we make progress towards this goal by obtaining near-optimal generators for two important special cases:

- We give a near optimal derandomization of the Chernoff bound for independent, uniformly random bits. Specifically, we show how to generate $x \in \{\pm 1\}^n$ using $\tilde{O}(\log(n/\varepsilon))$ random bits such that for any unit vector u, $u \cdot x$ matches the sub-Gaussian tail behaviour predicted by the Chernoff bound up to error ε.
- We construct a generator which fools halfspaces with $\{0, 1, -1\}$ coefficients with error ϵ with a seed-length of $\tilde{O}(\log(n/\epsilon))$. This includes the important special case of majorities.

In both cases, the best previous results required seed-length of $O(\log n + \log^2(1/\epsilon))$.

Technically, our work combines new Fourier-analytic tools with the iterative dimension reduction techniques and the gradually increasing independence paradigm of previous works [KMN11, CRSW13, GMR$^+$12].

1 Introduction

The theory of pseudorandomness has given compelling evidence that very strong pseudorandom generators (PRGs) exist. For example, assuming that there are computational problems solvable in exponential time that require exponential-sized circuits, Impagliazzo and Wigderson [IW97] showed that there exist very strong PRGs which allow us to simulate every randomized algorithm deterministically with only a polynomial slowdown, and thus BPP = P. These results, however, are conditional on a circuit complexity assumption whose proof seems far off. Since PRGs that fool a class of Boolean circuits also imply lower bounds for that class, we cannot hope to circumvent this assumption. Thus unconditional generators are only possible for restricted models of computation for which we have strong lower bounds.

Bounded-space algorithms are a natural computational model for which we know how to construct strong PRGs unconditionally. Let RL denote the class of randomized algorithms with $O(\log n)$ work space which can access the random bits in a read-once pre-specified order. Nisan [Nis92] devised a PRG of seed length $O(\log^2(n/\varepsilon))$ that fools RL. This generator was used by Nisan himself to show that RL $\subseteq$ SC [Nis94] and by Saks and Zhou [SZ99] to prove that RL can be simulated in space $O(\log^{3/2} n)$. Constructing PRGs with the optimal $O(\log(n/\epsilon))$ seed length for this class and showing that RL = L is arguably the outstanding open problem in derandomization (which might not require a breakthrough in lower bounds). Despite much progress in this area [INW94, NZ96, RR99, Rei08, RTV06, BRRY14, BV10, KNP11, De11, GMR$^+$12], there are few cases where we can improve on Nisan's twenty year old bound of $O(\log^2 n)$ [Nis92].

Halfspaces are Boolean functions $h : \{\pm 1\}^n \to \{\pm 1\}$ described as $h(x) = \mathrm{sgn}(\langle w, x\rangle - \theta)$ for some *weight* vector $w \in \mathbb{R}^n$ and *threshold* $\theta \in \mathbb{R}$. They are of central importance in computatonal complexity, learning theory and social choice. Lower bounds for halfspaces are trivial, whereas the problem of proving lower bounds against depth-2 TC_0 or halfspaces of halfspaces is a frontier open problem in computational complexity. The problem of constructing explicit PRGs that can fool halfspaces is a natural challenge that has seen a lot of exciting progress recently [DGJ$^+$09, MZ13, Kan11b, Kan14]. The best known PRG construction for halfspaces is that of Meka and Zuckerman [MZ13] who gave a PRG with seed-length $O(\log n + \log^2(1/\epsilon))$, which is $O(\log^2(n))$ for polynomially small error. They also made a connection to space bounded algorithms by showing that PRGs against RL with inverse polynomial error can be used to fool halfspaces. Thus constructing better PRGs for halfspaces seems to be a necessary step towards progress for bounded-space algorithms.

Beyond computational complexity, the problem of constructing better PRGs for halfspaces has ample algorithmic motivation; perhaps the most compelling of which comes from the ubiquitous applications in computer science of Chernoff-like bounds for weighted sums of the form $\sum_i w_i x_i$ where the x_is are uniformly random bits. There has been a long line of work on showing sharp tail bounds for pseudorandom sequences starting from [SSS95]. A PRG for halfspaces with seed-length $O(\log(n/\varepsilon))$ would give a space of support

size poly(n) where Chernoff-like tail bounds hold. This in turn would yield a black-box derandomization with only a polynomial slow-down of any algorithm which relies on uniform randomness only for such tail bounds. PRGs for halfspaces also have other algorithmic applications to streaming algorithms for duplicate detection [GR09] and efficient revenue maximization for certain kinds of auctions [GNR14].

1.1 Our results

A PRG is a function $\mathcal{G} : \{\pm 1\}^r \to \{\pm 1\}^n$. We refer to r as the seed-length of the generator. The $\tilde{O}()$ notation hides polylogarithmic factors in its argument. We say $\mathcal{G}$ is *explicit* if the output of $\mathcal{G}$ can be computed in time poly(n).

Definition 1. *A* PRG $\mathcal{G} : \{\pm 1\}^r \to \{\pm 1\}^n$ *fools a class of functions* $\mathcal{F} = \{f : \{\pm 1\}^n \to \{\pm 1\}\}$ *with error* ε *(or* ε*-fools* $\mathcal{F}$*) if for every* $f \in \mathcal{F}$*,*

$$\left| \Pr_{x \in \{\pm 1\}^n} [f(x) = 1] - \Pr_{y \in \{\pm 1\}^r} [f(G(y)) = 1] \right| < \varepsilon.$$

Derandomized Chernoff bounds

Chernoff bounds are a basic tool in the analysis of randomized algorithms. A ubiquitous version that applies to the setting of independent random bits is the following:

Claim 1 (Chernoff bound)**.** *There exist constants* $c_1, c_2 > 0$ *such that for every unit vector* $w \in \mathbb{R}^n$ *and* $t \geq 1$*,* $\Pr_{x \in \{\pm 1\}^n} [\, |\langle w,\ x \rangle| > t] \leq c_1 e^{-c_2 t^2}$*.*

We obtain a near-optimal derandomization of this result.

Theorem 1. *There exists an explicit generator* $\mathcal{G}_1 : \{\pm 1\}^r \to \{\pm 1\}^n$ *and constants* d_1, d_2 *such that for every unit vector* $w \in \mathbb{R}^n$*,* $t \geq 1$ *and* $\varepsilon > 0$*,*

$$\Pr_{y \in_u \{\pm 1\}^r} [\, |\langle w, \mathcal{G}_1(y) \rangle| > t] \leq d_1 e^{-d_2 t^2} + \varepsilon.$$

The generator has seed-length $r = \tilde{O}(\log(n/\varepsilon))$*.*

To contrast this with what was known previously, consider the setting where $\varepsilon = 1/\text{poly}(n)$. The Chernoff bound asserts that the probability that $|\langle w,\ x \rangle| = \Omega(\sqrt{\log(n)})$ is inverse polynomially small. A PRG for halfspaces with error parameter $\varepsilon = 1/\text{poly}(n)$ would also gurarantee such tails, but the best known construction requires seed-length $O(\log^2(n))$ [MZ13]. One could also get such tail bounds using limited indpendence [SSS95]; however, we would need $O(\log(n))$-wise independence, which again requires $O(\log^2(n))$ seed-length.

Fooling signed majorities

An important sub-class of halfspaces are those whose weight vectors have $\{0, 1, -1\}$-valued entries. This corresponds to selecting a subset of variables, assiging each of them an orientation and then taking a threshold. We henceforth refer to this class of halfspaces as signed majorities. Signed majorities arise naturally in voting theory, learning theory and property testing - see [MOO05, MORS09, RS13, BO10]. Fooling such tests requires fooling the sum of arbitrary subsets of variables in statistical distance, a problem that was studied by [GMRZ13] in their work on fooling *combinatorial shapes*. Fooling sums in statistical distance includes as a special case modular tests on sums of variables with unrestricted modulus [LRTV09, MZ09]. PRGs for modular sums are a strong generalization of the versatile small-bias spaces [NN93] which correspond to fooling modular sums with modulus two. The best previously known PRGs due to Lovett *et al.* for such tests require seed-length $O(\log^2 n)$ [LRTV09] for large modulii, but their result can also handle sums with non-binary coefficients. Finally, signed majorities seem to capture several technical hurdles in designing optimal PRGs for halfspaces.

We construct a PRG which ε-fools signed majorities with a seed-length of $\tilde{O}(\log(n/\varepsilon))$.

Theorem 2. *There exists an explicit generator $\mathcal{G}_2 : \{\pm 1\}^r \to \{\pm 1\}^n$ with seed-length $r = \tilde{O}(\log(n/\varepsilon))$ which ε-fools signed majorities .*

The best previous result even for signed majorities had a seed-length of $O(\log n + \log^2(1/\varepsilon))$ [MZ13]. For the important case of polynomially small error, $\varepsilon = 1/\mathrm{poly}(n)$, our result gives the first improvement over the $O(\log^2 n)$ bound implied by directly applying known PRGs for space-bounded machines [Nis92, INW94].

Independently and concurrently De [De14] gave a PRG for *combinatorial shapes* introduced by [GMRZ13] with a seed-length of $O(\log^{3/2}(n/\varepsilon))$. These objects are more general than signed majorities but De's seed-length is worse than ours.

1.2 Other related work

Starting with the work of Diakonikolas et al. [DGJ$^+$09], there has been a lot of work on constructing PRGs for halfspaces and related classes of intersections of halfspaces and polynomial threshold functions over the domain $\{\pm 1\}^n$ [DKN10, GOWZ10, HKM12, MZ13, Kan11b, Kan11a, Kan14]. Rabani and Shpilka [RS10] construct optimal hitting set generators for halfspaces over $\{\pm 1\}^n$; hitting set generators are in general weaker than PRGs however.

Another line of work gives constructions of PRGs for halfspaces for the uniform distribution over the sphere (*spherical caps*) or the Gaussian distribution. This case is easier than constructing PRGs for halfspaces over the hypercube; the latter objects are known to imply the former with comparable parameters. For spherical caps, Karnin, Rabani and Shpilka [KRS12] gave a PRG with a seed-length of $O(\log n + \log^2(1/\varepsilon))$. For the Gaussian distribution, [Kan14] gave a PRG which achieves a seed-length of $O(\log n + \log^{3/2}(1/\varepsilon))$. Very

recently, Kothari and Meka [KM14] gave a PRG for spherical caps with a seed-length of $\tilde{O}(\log(n/\varepsilon))$. At a high level, [KM14] also uses the iterative dimension reduction approach like in [KMN11, CRSW13, GMR$^+$12]; however, the final construction and its analysis are significantly different.

1.3 Overview of our constructions

Derandomized Chernoff bounds

Our first attempt at constructing a PRG for the Chernoff bound applies a simple dimension reduction step iteratively.

1. Starting from a linear function $\sum_{i=1}^{n} w_i x_i$, (pseudo)randomly hash the variables into $\sqrt{n}$ buckets using a hash function h.
2. Use an ε-biased string x to sum up coefficients within a bucket. This gives a new linear function $\sum_{j=1}^{\sqrt{n}} v_j y_j$ in $\sqrt{n}$ dimensions where $v_j = \sum_{i:h(i)=j} w_i x_i$.

Repeating this step $\log\log(n)$ times, we get down to $\theta \in \mathbb{R}$ which is the value we output. Call this generator $\mathcal{G}'$. It is easy to see that each output bit of $\mathcal{G}'$ is the xor of $\log\log(n)$ bits from independent ε-biased strings, where the hash functions are used to select co-ordinates from each string. This technique of applying pseudorandom dimension reduction iteratively is similar to [KMN11, CRSW13, GMR$^+$12].

Does this generator give the desired tail behavior? Assume that we start from a unit vector $w \in \mathbb{R}^n$. To get tail bounds, we would like to control the ℓ_2 norm, which starts at 1 but could increase substantially for particular choices of x. The Chernoff bound says that for truly random x, the ℓ_2 norm is unlikely to increase by more than a factor of $c\sqrt{\log(n)}$. Even if we manage to match this tail behavior in each step by choosing x pseudorandomly (which is the problem we are trying to solve), the final bound we get would be $O((\log n)^{\log\log(n)/2})$. Using ε-biased x, we show a weaker bound of polylog(n) for each step, giving an overall bound of $d(n) = (\log(n))^{O(\log\log(n))}$. Showing this bound for one step requires a fair amount of technical work, it works by decomposing the vector into weight scales and tuning the amount of independence to the scale like in [GMR$^+$12]. We leave open the question of whther $\mathcal{G}'$ can itself give Chernoff-like tail bounds.

Next we show that one gets the desired tail behaviour by hashing variables into $m = \text{poly}(d(n))$ buckets and using an independent copy of $\mathcal{G}'$ for each bucket. The reason is the output of the resulting generator can be viewed as the sum of m independent *bounded* random variables, which lets us apply Bernstein's inequality which guarantees Chernoff-like tails for such variables. The boundedness comes from the tail guarantee of $\mathcal{G}'$: since large deviations are very unlikely, we can condition on the event that they do not occur in any of the buckets. The final step is to reduce to seed-length, we do this by recycling the seed for the various independent copies of $\mathcal{G}'$ using the INW generator [INW94], like in [MZ13].

Fooling signed majorities

Let us fix a test vector $v \in \{-1, 0, 1\}^n$ and error $\varepsilon = 1/\mathrm{poly}(n)$. Fooling signed majorities with polynomially small error is equivalent to fooling *linear sums* of the form $\langle v, x\rangle$ in statistical distance with error $1/\mathrm{poly}(n)$. We shall adopt this view from now on.

We start with a generator that uses iterated dimension reduction and gradually-increasing independence as we did for derandomizing the Chernoff bound. This by itself is not enough for fooling sums in statistical distance. The reason is that there exist small-bias spaces with exponentially small bias that are far from fooling linear sums in statistical distance, like the set of strings whose weight is divisible by 3 [VW08]. We design a different generator to deal with such tests and then combine the two generators by xoring independent copies.

Next, note that showing closeness in statistical distance for discrete random variables is equivalent to showing that their *Fourier transforms* are close. Using this, it suffices to design a generator $\mathcal{G} : \{\pm 1\}^r \to \{\pm 1\}^n$ such that for all $\alpha \in \mathbb{R}$, the corresponding *Fourier coefficient* $\mathbb{E}_y[\exp(2\pi i \alpha \langle v, \mathcal{G}(y)\rangle)]$ is close to its value under the uniform distribution. Note that in order to fool the mod m test, it suffices to fool all $\alpha = j/m$ for integers j. We consider two cases based on how large α is relative to $\|v\|_0 = k$.

Large α: Here we consider $\alpha \gg 1/\sqrt{k}$. This includes the case of modular tests where the modulus is much smaller than $\sqrt{k}$. We fool such tests using an error reduction procedure. We start with the generator of [GMRZ13] which requires seed-length $O(\log n)$ to fool such tests with constant error. We then reduce the error to inverse polynomial at the expense of a $O(\log \log n)$ factor in seed-length using standard machinery from pseudorandomness. While technically simple, this step is the botteleneck in extending our result to more general halfpsaces: there is no analog of the [GMRZ13] generator to start from.

Smaller α: This case which includes modular tests where the modulus is $\Omega(\sqrt{k})$ is the harder case and technically the most novel portion of this work. The qualitative difference from the other case can be seen from the fact that when we sum k random bits modulo $m = \omega(\sqrt{k})$, the resulting distribution is no longer uniform over congruence classes.

The generator uses dimension reduction in a manner similar to what we used to derandomize the Chernoff bound. Like before, the plan is to show that a single dimension reduction step does not incur too much error. However, the analysis is very different and requires several new tools. This step critically exploits the recursive structure of the generator: to analyze the error we can work as if the variables in the reduced space are given truly random signs and then recursively analyze the error in the reduced space. Working with truly random bits in the smaller-dimensional space helps us reduce bounding the error to finding good low-degree polynomial approximators for a certain product of cosines. In the most technically, involved part of our argument we use various analytic tools to find such low-degree approximators. One additional ingredient is that the above approach does not actually work for all test vectors but only for *sufficiently well-spread out* vectors as

measured by their ℓ_2, ℓ_4 norms. The final piece is to argue that the ℓ_2, ℓ_4 norms are not distorted too much by the dimension reduction steps.

2 Preliminaries

We start with some notation:

- For vectors $x \in \mathbb{R}^n$, let $\|x\|_p$ denote the usual ℓ_p-norms, and let $\|x\|_0$ denote the size of the support of x. For a random variable X and $p > 0$, let $\|X\|_p = \mathbb{E}[|X|^p]^{1/p}$.
- For a multi-linear polynomial $Q : \mathbb{R}^n \to \mathbb{R}$, $\|Q\|_2^2$ denotes the sum of squares of coefficients of Q and $\|Q\|_1$ denotes the sum of absolute values of the coefficients.
- For vectors $u, v \in \mathbb{R}^n$, let $u \star v = (u_i v_i)_{i=1}^n$ denote the coordinate-wise product.
- For $v \in \{\pm 1\}^n$ and $\alpha \in \mathbb{R}$, define $\phi_{v,\alpha}(x) = \exp(2\pi i \alpha (v \cdot x))$.
- For $v \in \mathbb{R}^n$ and a hash function $h : [n] \to [m]$, define

$$h(v) = \sum_{j=1}^{m} \|v_{|h^{-1}(j)}\|_2^4 \tag{1}$$

- For a hash function $h : [n] \to [m]$, let $A(h) \in \{0,1\}^{m \times n}$, be the matrix with $A(h)_{ji} = 1$ if and only if $h(i) = j$.
- For a string $x \in \{\pm 1\}^n$, let $D(x) \in \mathbb{R}^{n \times n}$ be the diagonal matrix formed by x.
- For two random variables X, Y over a domain Ω, their statistical distance is defined as $d_{TV}(X, Y) = \max_{A \subseteq \Omega} |\Pr[X \in A] - \Pr[Y \in A]|$.
- Unless otherwise stated c, C denote universal constants.

Throughout we assume that n is sufficiently large and that $\delta, \epsilon > 0$ are sufficiently small. All vectors here will be row vectors rather than column vectors.

Definition. *For $n, m, \delta > 0$ we say that a family of hash functions $\mathcal{H} = \{h : [n] \to [m]\}$ is δ-biased if for any $r \leq n$ distinct indices $i_1, i_2, \ldots, i_r \in [n]$ and $j_1, \ldots, j_r \in [m]$,*

$$\Pr_{h \in_u \mathcal{H}}[h(i_1) = j_1 \wedge h(i_2) = j_2 \wedge \cdots \wedge h(i_r) = j_r] = \frac{1}{m^r} \pm \delta.$$

We say that such a family is k-wise independent if the above holds with $\delta = 0$ for all $r \leq k$. We say that a distribution over $\{\pm 1\}^n$ is δ-biased or k-wise independent if the corresponding family of functions $h : [n] \to [2]$ is.

Such families of functions can be generated using small seeds.

Fact 3. *For $n, m, k, \delta > 0$, there exist explicit δ-biased families of hash functions $h : [n] \to [m]$ that are generated from a seed of length $s = O(\log(n/\delta))$. There are also, explicit k-wise independent families that are generated from a seed of length $s = O(k \log(nm))$.*

Taking the pointwise sum of such generators modulo m gives a family of hash functions that is both δ-biased and k-wise independent generated from a seed of length $s = O(\log(n/\delta) + k \log(nm))$.

2.1 Basic Results

We collect some known results about pseudorandomness and prove some other technical results that will be used later. We defer all proofs for this section to Appendix A.

We will use the following result from [GMRZ13] giving PRGs for signed majorities.

Theorem 4. *[GMRZ13] For $n, \varepsilon > 0$ there exists an explicit pseudorandom generator, $Y \in \{\pm 1\}^n$ generated from a seed of length $s = O(\log(n) + \log^2(1/\epsilon))$ so that for any $v \in \{-1, 0, 1\}^n$ and $X \in_u \{\pm 1\}^n$, we have that $d_{TV}(v \cdot Y, v \cdot X) \leq \epsilon$.*

We shall use PRGs for small-space machines or read-once branching programs (ROBP) of Nisan [Nis92] and Impagliazzo, Nisan and Wigderson [INW94].

Definition 2 ((S, D, T)-ROBP)**.** *An (S, D, T)-ROBP M is a layered directed graph with $T + 1$ layers and 2^S vertices per layer with the following properties.*

- *The first layer has a single* start *node and the last layer has two nodes labeled $0, 1$ respectively.*
- *A vertex v in layer i, $0 \leq i < T$ has 2^D edges to layer $i + 1$ each labeled with an element of $\{0, 1\}^D$.*

A graph M as above naturally defines a function $M : \left(\{0,1\}^D\right)^T \to \{0, 1\}$ where on input $(z_1, \ldots, z_T) \in \left(\{0,1\}^D\right)^T$ one traverses the edges of the graph according to the labels $z_1, \ldots, z_T$ and outputs the label of the final vertex reached.

Theorem 5 ([Nis92], [INW94])**.** *There exists an explicit PRG $\mathcal{G}^{INW} : \{0,1\}^r \to \left(\{0,1\}^D\right)^T$ which ϵ-fools (S, D, T)-branching programs and has seed-length $r = O(D + S \log T + \log(T/\delta) \cdot (\log T))$.*

We will need to make use of the hypercontractive inequality (see [O'D14]):

Lemma 6 (Hypercontractivity)**.** *Let $x \sim_u \{\pm 1\}^n$. Then, for a degree d polynomial Q and an even integer $p \geq 2$,*

$$\mathbb{E}\left[Q(x)^p\right] \leq (p-1)^{pd/2} \cdot \|Q\|_2^p.$$

Lemma 7 (Hypercontractivity δ-biased)**.** *Let $x \sim \mathcal{D}$ be drawn from a δ-biased distribution. Then, for a degree d polynomial Q and an even integer $p \geq 2$,*

$$\mathbb{E}\left[Q(x)^p\right] \leq (p-1)^{pd/2} \cdot \|Q\|_2^p + \|Q\|_1^p \delta.$$

We will use the following Chernoff-like tail bound for small-bias spaces.

Lemma 8. *For all $v \in \mathbb{R}^n$ with $\|v\|_2 = 1$ and $x \sim \mathcal{D}$ ϵ-biased over $\{\pm 1\}^n$, and $t \geq 1$*

$$\Pr\left[|\langle v,\ x\rangle| > t\right] \leq 2\exp(-t^2/4) + \|v\|_0^{t^2} \cdot \epsilon.$$

The next two lemmas quantify load-balancing properties of δ-biased hash functions in terms of the ℓ_p-norms of vectors.

Lemma 9. *Let $p \geq 2$ be an integer. Let $v \in \mathbb{R}^n$ and $\mathcal{H} = \{h : [n] \to [m]\}$ be either a δ-biased hash family for $\delta > 0$ or a p-wise independent family for $\delta = 0$. Then*

$$\mathbb{E}[h(v)^p] \leq O(p)^{2p} \left(\frac{\|v\|_2^4}{m}\right)^p + O(p)^{2p}\|v\|_4^{4p} + m^p \|v\|_2^{4p}\delta.$$

Lemma 10. *For all $v \in \mathbb{R}_+^n$ and $\mathcal{H} = \{h : [n] \to [m]\}$ a δ-biased family, and $j \in [m]$, and all even $p \geq 2$,*

$$\Pr\left[\left|\left\|v_{|h^{-1}(j)}\right\|_1 - \|v\|_1/m\right| \geq t\right] \leq \frac{O(p)^{p/2}\,\|v\|_2^p + \|v\|_1^p\,\delta}{t^p}.$$

3 Derandomizing the Chernoff Bounds

In this section we present a pseudorandom generator that gives Chernoff-like tail bounds.

Theorem 11. *For all $\delta > 0$, there exists an explicit generator $\mathcal{G} : \{0,1\}^r \to \{\pm 1\}^n$ with seed-length $r = \tilde{O}(\log(n/\delta))$ such that for all unit vectors $w \in \mathbb{R}^n$, and $t \geq 0$,*

$$\Pr_{y \in_u \{0,1\}^r}[|\langle w,\ \mathcal{G}(y)\rangle| \geq t] \leq 4\exp(-t^2/16) + \delta.$$

Our construction proceeds in two steps. We first construct a generator which has moderate tail bounds but does not match the tail behaviour of truly random distribution. We then boost the tail bounds to match the behaviour of truly random distributions using PRGs for small-space machines.

3.1 Moderately Decaying Tails

The main result of this section is a generator with the following tail behaviour.

Lemma 12. *For n and $\gamma \in (0,1)$, there exists an explicit generator $\mathcal{G}' : \{0,1\}^{r'} \to \{\pm 1\}^n$ with seed-length $r' = O(\log(n/\gamma)\log\log(n))$ such that for all unit vectors $w \in \mathbb{R}^n$,*

$$\Pr_{y \in_u \{0,1\}^r}\left[|\langle w,\ \mathcal{G}'(y)\rangle| \geq (C_1 \log(n/\gamma))^{C_2 \log\log(n)}\right] \leq \gamma.$$

The generator $\mathcal{G}'$ is recursively defined. We first specify the *one-step* generator $\mathcal{G}''$ that is used in defining $\mathcal{G}'$. Fix $\delta > 0$, n. Let $\mathcal{H} = \{h : [n] \to [m]\}$ be a family of δ-biased hash functions. Let $\mathcal{D}$ be a δ-biased distribution over $\{\pm 1\}^n$. The generator $\mathcal{G}''$ takes as input a hash function $h \in \mathcal{H}$, $x \in \mathcal{D}$ and $z \in_u \{\pm 1\}^m$, the output is

$$\mathcal{G}''(h,x,z) = zA(h)D(x).$$

Thus we have for any $i \in [n]$,

$$\mathcal{G}''(h,x,z)_j = z_{h(i)}x_i.$$

Thus the generator $\mathcal{G}''$ starts with the δ-biased string $x \in \mathcal{D}$ as output, hashes the coordinates into $[m]$ bins and flips the signs of all coordinates in each bin by picking a uniformly random independent bit for each bin. This takes $O(\log(n/\delta) + m)$ random bits.

The generator $\mathcal{G}'$ is obtained by taking $m \approx \sqrt{n}$ and then recursively using $\mathcal{G}''$ to generate $z \in \{\pm 1\}^n$. The base case of the recursion is reached when $m = O(\log(n/\delta))$ at which point we use a truly random string z. This requires $k \leq \log\log(n)$ stages of recursion, so that the seed length is $O(\log(n/\delta)\log\log(n))$. Unrolling the recursion, we see that if we set $n_\ell = n^{2^{-\ell}}$ for $\ell \in \{0,\ldots,k\}$ then $\mathcal{G}'$ takes as input two sequences:

- A sequence of hash functions $h^1,\ldots,h^k$ where $h^\ell : [n_{\ell-1}] \to [n_\ell]$ is drawn from a δ-biased family of hash functions.
- A sequence of strings $x^1,\ldots,x^k$ where $x^\ell \in \{\pm 1\}^{n_\ell}$ is drawn from a δ-biased distribution.

For each coordinate i, consider the sequence $\{i_\ell \in n_\ell\}_{\ell=0}^k$ obtained by successively applying the hash functions:

$$i_0 = i, \quad i_\ell = h^\ell(i_{\ell-1}) \text{ for } \ell \geq 1.$$

Then we have

$$\mathcal{G}'(h^1,\ldots,h^k,x^1,\ldots,x^k) = \prod_{\ell=1}^k x^\ell_{i_\ell}.$$

The analysis of $\mathcal{G}'$ proceeds step by step and each step reduces to analyzing $\mathcal{G}''$. Note that

$$\langle w,\ \mathcal{G}''(h,x,z)\rangle = \langle w,\ zA(h)D(x)\rangle = \langle A(h)D(x)w^T,\ z\rangle.$$

We can view $A(h)D(x)w^T$ as projection of $w \in \mathbb{R}^n$ down to $\mathbb{R}^m$ where we first hash coordinates into buckets, and then sum the coordinates in a bucket with signs given by x. The next lemma saying that the transformation $A(h)D(x)$ is unlikely to stretch Euclidean norms too much serves as the base case for the recursion.

Lemma 13. *Let $n \geq 1$ and $m = \sqrt{n}$ and $\delta < 1/10n^2$. Let $\mathcal{D}$ be a δ-biased distribution over $\{\pm 1\}^n$ and $\mathcal{H} = \{h : [n] \to [m]\}$ be a δ-biased hash family. There exists a constant C such that for all unit vectors $w \in \mathbb{R}^n$,*

$$\Pr_{x \in \mathcal{D}, h \in \mathcal{H}} \left[\left\| A(h)D(x)w^T \right\|_2 \geq C(\log\log(n)) \log(1/\delta)^{3/4} \right] \leq 3(\log\log(n))m\sqrt{\delta}$$

We prove this lemma by decomposing the vector w across various weight scales. Fix a unit vector $w \in \mathbb{R}^n$. Without loss of generality, we ignore all coordinates i where $|w_i| \leq 1/n$ as they can only effect the ℓ_2-norm by at most 1. For $\ell \in \{1, \ldots, \log\log n\}$, define $w(\ell) \in \mathbb{R}^n$ as

$$w(\ell)_i = \begin{cases} w_i & \text{if } |w_i| \in \left(\frac{1}{2^{2^\ell}}, \frac{1}{2^{2^{\ell-1}}} \right] \\ 0 & \text{otherwise.} \end{cases}$$

Thus $w(\ell)$ picks out the entries in the ℓ^{th} weight scale. In addition, define $w(0)$ to consist of entries that lie in the interval $(1/2, 1]$. We will show that for every ℓ, the bound

$$\left\| A(h)D(x)w(\ell)^T \right\|_2^2 \leq O(1) \log(1/\delta)^{1.5}$$

holds with high (inverse polynomial) probability. Here we tailor the amount of independence we use in the argument to the scale, in a manner similar to [CRSW13, GMR+12]. Once this is done, Lemma 13 follows by the triangle inequality.

We start with a simple bound which suffices for small constant ℓ.

Lemma 14. *For all ℓ, we have*

$$\left\| A(h)D(x)w(\ell)^T \right\|_2 \leq 2^{2^\ell}.$$

Proof. Observe that

$$\left\| A(h)D(x)w(\ell)^T \right\|_\infty \leq \|w(\ell)\|_1$$
$$\left\| A(h)D(x)w(\ell)^T \right\|_1 \leq \|w(\ell)\|_1$$

hence by Holder's inequality

$$\left\| A(h)D(x)w(\ell)^T \right\|_2 \leq \|w(\ell)\|_1 .$$

Since $\|w(\ell)\|_2 \leq 1$ and every non-zero entry is at least 2^{-2^ℓ} we have $\|w(\ell)\|_1 \leq 2^{2^\ell}$ and hence

$$\left\| A(h)D(x)w(\ell)^T \right\|_2 \leq 2^{2^\ell}.$$

□

Given this lemma, we can assume that ℓ is a sufficiently large constant. We show that the weight vector is hashed fairly regularly with high probability over the choice of $h \in \mathcal{H}$, where the regularity is measured by $h(v)$.

Lemma 15. *Fix $\ell \geq 2$. Then*

$$\Pr_{h\in\mathcal{H}}\left[h(w(\ell)) \leq 2C_4\frac{\sqrt{\log(1/\delta)}}{2^{2^{\ell-2}}}\right] \geq 1 - 2m\sqrt{\delta}. \tag{2}$$

Proof. By Lemma 10 applied to the vector $(w(\ell)_i^2)_{i=1}^n$, there is a constant C_4 such that for even $q \geq 2$

$$\Pr_{h\in\mathcal{H}}\left[h(w(\ell)) \geq 1/m + t\right] \leq m\left(\left(\frac{C_4\sqrt{q}\,\|w(\ell)\|_4^2}{t}\right)^q + \delta\left(\frac{\|w(\ell)\|_2^2}{t}\right)^q\right).$$

Plugging in the bounds

$$\|w(\ell)\|_4^2 \leq \|w(\ell)\|_\infty \leq 2^{-2^{\ell-1}},\ \|w(\ell)\|_2 \leq 1$$

we get

$$\Pr_{h\in\mathcal{H}}\left[h(w(\ell)) \geq 1/m + t\right] \leq m\left(\left(\frac{C_4\sqrt{q}}{2^{2^{\ell-1}}t}\right)^q + \delta\left(\frac{1}{t}\right)^q\right).$$

Therefore, taking

$$q = \frac{\log(1/\delta)}{2^{\ell-1}},\quad t = \frac{C_4\sqrt{q}}{2^{2^{\ell-2}}}$$

in the above equation, we get

$$\begin{aligned}\frac{C_4\sqrt{q}\,\|w(\ell)\|_4^2}{t} &\leq \frac{2^{2^{\ell-2}}}{2^{2^{\ell-1}}} \leq \frac{1}{2^{2^{\ell-2}}},\\ \left(\frac{C_4\sqrt{q}\,\|w(\ell)\|_4^2}{t}\right)^q &\leq \frac{1}{2^{2^{\ell-2}\log(1/\delta)/2^{\ell-1}}} \leq \sqrt{\delta}.\\ \left(\frac{1}{t}\right)^q &\leq (2^{2^{\ell-2}})^{\log(1/\delta)/2^{\ell-1}} = \frac{1}{\sqrt{\delta}}.\end{aligned}$$

hence

$$\Pr_{h\in\mathcal{H}}\left[h(w(\ell)) \geq 1/m + C_4\frac{\sqrt{\log(1/\delta)}}{2^{2^{\ell-2}}}\right] \leq 2m\sqrt{\delta}.$$

Hence with probability $1 - 2m\sqrt{\delta}$ over the choice of h, we have

$$h(w(\ell)) \leq \frac{1}{m} + \frac{C_4\sqrt{\log(1/\delta)}}{2^{2^{\ell-2}}} \leq 2C_4\frac{\sqrt{\log(1/\delta)}}{2^{2^{\ell-2}}}$$

since $1/m = 1/\sqrt{n} \leq 2^{-2^{\ell-2}}$. ∎

Conditioned on the hash function h being good (i.e., satisfying the condition of the previous lemma), we will show that $\left\|A(h)D(x)w(\ell)^T\right\|_2^2$ is small with high-probability.

Lemma 16. *Fix $\ell \geq 2$ and assume that h is such that the event described in Equation (2) holds. There exists a constant C_6 such that*

$$\Pr_{x\in\mathcal{D}}\left[\left\|A(h)D(x)w(\ell)^T\right\|_2 \leq C_6 \log\left(\frac{1}{\delta}\right)^{5/8}\right] \geq 1-\sqrt{\delta}-\delta^{16}. \tag{3}$$

Proof. We will show that $\left\|A(h)D(x)w(\ell)^T\right\|_2^2$ is concentrated around its mean (which is $\|w(\ell)\|_2^2$) by bounding its moments. The deviation is given by the polynomial

$$\begin{aligned}
Q_\ell(x) &= \left\|A(h)D(x)w(\ell)^T\right\|_2^2 - \|w(\ell)\|_2^2 \\
&= \sum_{j\in[m]}\left(\sum_{i\in h^{-1}(j)} w(\ell)_i x_i\right)^2 - \|w(\ell)\|_2^2 \\
&= \sum_{j\in[m]}\sum_{i_1\neq i_2\in h^{-1}(j)} w(\ell)_{i_1} w(\ell)_{i_2} x_{i_1} x_{i_2}
\end{aligned}$$

We have

$$\begin{aligned}
\mathop{\mathbb{E}}_{x\in_u\{\pm1\}^n}[Q_\ell(x)^2] &= \sum_{j=1}^{m}\sum_{i_1\neq i_2\in h^{-1}(j)} (w(\ell)_{i_1})^2 (w(\ell)_{i_2})^2 \\
&\leq \sum_{j=1}^{m} \|w(\ell)_{|h^{-1}(j)}\|_2^4 \\
&= h(w(\ell)).
\end{aligned}$$

$$\begin{aligned}
\|Q_\ell\|_1 &= \sum_{j\in[m]}\sum_{i_1\neq i_2\in h^{-1}(j)} |w(\ell)_{i_1} w(\ell)_{i_2}| \\
&\leq \sum_{j\in m} \left\|w(\ell)_{|h^{-1}(j)}\right\|_1^2 \\
&\leq \|w(\ell)\|_1^2 \\
&\leq \|w(\ell)\|_0 .
\end{aligned}$$

By Lemma 7 applied to Q with $d=2$, there exists a constant C_3 so that for all even $p \geq 2$,

$$\mathop{\mathbb{E}}_{x\in\mathcal{D}}[Q(x)^p] \leq \left(C_3 p\sqrt{h(w(\ell))}\right)^p + \|w(\ell)\|_0^p\,\delta. \tag{4}$$

We bound $h(w_\ell)$ using Equation (2). We also have $\|w(\ell)\|_0 \leq 2^{2^\ell}$. Plugging these into Equation (4),

$$\mathop{\mathbb{E}}_{x\in\mathcal{D}}[Q(x)^p] \leq \frac{C_5^p p^p \log(1/\delta)^{p/4}}{2^{2^{\ell-3}p}} + 2^{2^\ell p}\delta$$

Now setting

$$p = \frac{\log(1/\delta)}{2^{\ell+1}},\ \theta = C_5 \log(1/\delta)^{5/4}$$

and using Markov' inequality gives

$$\Pr_{x \in \mathcal{D}}[Q(x) \geq \theta] \leq \left(\frac{C_5 p \log(1/\delta)^{1/4}}{2^{2^{\ell-3}}\theta} \right)^p + \left(\frac{2^{2^\ell}}{\theta} \right)^p \delta$$

To bound the first term, note that

$$\left(\frac{C_5 p \log(1/\delta)^{1/4}}{2^{2^{\ell-3}}\theta} \right)^p \leq \left(\frac{1}{2^{2^{\ell-3}}} \right)^{\log(1/\delta)/2^{\ell+1}} \leq \delta^{16}.$$

For the second term, note that since $\theta \geq 1$,

$$\left(\frac{2^{2^\ell}}{\theta} \right)^p \delta \leq 2^{2^\ell \log(1/\delta)/2^{\ell+1}} \delta \leq \sqrt{\delta}$$

Therefore, except with probability at least $\sqrt{\delta} + \delta^{16}$ we have

$$\left\| A(h) D(x) w(\ell)^T \right\|_2^2 \leq \|w_\ell\|_2^2 + C_5 \log(1/\delta)^{5/4}$$

hence

$$\left\| A(h) D(x) w(\ell)^T \right\|_2 \leq C_6 \log(1/\delta)^{5/8}$$

□

We now finish the proof of Lemma 13.

Proof of Lemma 13. Note that

$$w = \sum_{i=0}^{\log\log(n)} w(\ell).$$

We will assume that h and x are chosen so that the conditions in Equations (2) and (3) hold for all ℓ. By the union bound, this happens except with probability

$$\log\log(n)(2m\sqrt{\delta} + \sqrt{\delta} + \delta^{16}) < 3\log\log(n) m\sqrt{\delta}.$$

In which case, we have

$$\begin{aligned}\|A(h)D(x)w\|_2 &= \left\| \sum_{\ell=0}^{\log\log(n)} A(h)D(x)w(\ell)^T \right\|_2 \\ &\leq \sum_{\ell=0}^{\log\log(n)} \left\|A(h)D(x)w(\ell)^T\right\|_2 \\ &\leq C_6 \log\log(n) \log(1/\delta)^{5/8} \\ &\leq C_6 \log(1/\delta)^{3/4}.\end{aligned}$$

□

We now prove the main lemma of this section:

Proof of Lemma 12. Let $k = \log\log(n)$ be the number of recursive stages. Let w be a unit vector. Given $h^1, \ldots, h^k$ and $x^1, \ldots, x^k$, we have

$$\langle w,\ \mathcal{G}'(h^1, \ldots, h^k, x^1, \ldots, x^k)\rangle = \prod_{\ell=1}^{k} A(h^\ell)D(x^\ell)w^T$$

Let $C_8 n\sqrt{\delta} = \gamma$, so that $\delta = \Omega(\gamma^2/n^2)$. Note that $\log(1/\delta) \gg \log\log(n)^4$.

By applying Lemma 13 inductively and using the union bound, except with probability

$$3(\log\log(n))^2 m\sqrt{\delta} \leq C_8 n\sqrt{\delta}$$

we have that for every $i \leq k$

$$\left\| \prod_{\ell=1}^{i} A(h^\ell)D(x^\ell)w^T \right\|_2 \leq (C_6 \log\log(n) \log(1/\delta)^{3/4})^i \leq \log(1/\delta)^{i/2}$$

and hence

$$|\langle w,\ \mathcal{G}'(h^1, \ldots, h^k, x^1, \ldots, x^k)\rangle| \leq \log(1/\delta)^{\log\log(n)/2}.$$

Thus, with probability $1 - \gamma$, the deviation is bounded by

$$d(n, \gamma) := (C_1 \log(n/\gamma))^{C_2 \log\log(n)}$$

and the seedlength of this generator is

$$r' = O(\log(n/\delta)\log\log(n)) = O(\log(n/\gamma)\log\log(n)).$$

□

3.2 Getting sub-Gaussian tail bounds

The generator $\mathcal{G}'$ gives a tail probability of $1-\gamma$ pseudorandomly for $d(n,\gamma)$ standard deviations. We now boost this to obtain sub-Gaussian tails by starting with independent copies of $\mathcal{G}'$ and then reuse the seeds for using a PRG for space bounded computations.

We will make some added assumptions about $\mathcal{G}'$:

- The output is ε-biased for some $\varepsilon \ll \gamma$. We can ensure this by xor-ing the output with an ε-biased string.
- The distribution is symmetric: for every x, $\Pr_y[\mathcal{G}'(y) = x] = \Pr_y[\mathcal{G}'(y) = -x]$. We ensure this by outputting either $\mathcal{G}'(y)$ or $-\mathcal{G}'(y)$ with probability $1/2$.

Let D_1, D_2 be constants such that

$$m = (D_1 \log(n/\gamma))^{D_2 \log\log(n)} > 10 d(n,\gamma)^2 \log(1/\gamma).$$

Note that for $\gamma = 1/\mathrm{poly}(n)$, $\log(m) = O(\log\log(n)^2)$.

Let $\mathcal{H} = \{h : [n] \to [m]\}$ be a family of γ-biased hash functions. Define a new generator $\bar{\mathcal{G}} : (\{0,1\}^r)^m \times \mathcal{H} \to \{\pm 1\}^n$ as follows:

$$\bar{\mathcal{G}}(z_1, \ldots, z_T, h)_i = \mathcal{G}'(z_j)_i, \text{ if } h(i) = j. \tag{5}$$

The seed-length of the generator is $\bar{r} = \log(n/\delta) + m \cdot r'$ which we will later improve to $\log(n/\delta) + r' + \log(n/\gamma)\log(m)$ using PRGs for space bounded computations.

The following claim characterizes the tail behaviour of the output of $\bar{\mathcal{G}}$.

Lemma 17. *Let $0 < \varepsilon < \gamma \le 1/n^3$. For all unit vectors $w \in \mathbb{R}^n$, the generator $\bar{\mathcal{G}}$ satisfies*

$$\Pr_{y \in \{0,1\}^{\bar{r}}}\left[\left|\langle w,\ \bar{\mathcal{G}}(y)\rangle\right| \ge t\right] \le 4(\exp(-t^2/16) + m\sqrt{\gamma} + (1/\gamma)^{4\log(m)}\varepsilon).$$

Proof. Note that it suffices to prove the claim for $t \le 2\sqrt{\log(1/\gamma)}$ since tail probabilities can only decrease with t and beyond this value, the tail bound is dominated by the additive terms.

Fix a unit vector $w \in \mathbb{R}^n$. Let

$$\beta = \frac{1}{m^2\sqrt{\log(1/\gamma)}} \tag{6}$$

and define $u, v \in \mathbb{R}^n$ to consist of the heavy and light indices respectively

$$u_i = \begin{cases} w_i & \text{if } |w_i| \ge \beta \\ 0 & \text{otherwise.} \end{cases}, \quad v_i = \begin{cases} w_i & \text{if } |w_i| < \beta \\ 0 & \text{otherwise.} \end{cases}.$$

Since $w = u + v$, it suffices to bound the probability that either of $|\langle u,\ \bar{\mathcal{G}}(y)\rangle|$ and $|\langle v,\ \bar{\mathcal{G}}y\rangle|$ exceed $t/2$. We will consider u first. Note that

$$\|u\|_0 \leq \frac{1}{\beta^2} \leq m^4 \log\left(\frac{1}{\gamma}\right).$$

Since $\bar{\mathcal{G}}(y)$ is ε-biased, by Lemma 8 applied for $t/2 \leq \sqrt{\log(1/\gamma)}$,

$$\begin{aligned}\Pr\left[\left|\langle u,\ \bar{\mathcal{G}}(y)\rangle\right| > t/2\right] &\leq 2\exp(-t^2/16) + \|u\|_0^{\log(1/\gamma)}\,\varepsilon \\ &\leq 2\exp(-t^2/16) + (m\log(1/\gamma))^{\log(1/\gamma)}\varepsilon. \end{aligned} \tag{7}$$

We will show a tail bound for $\langle v,\ \bar{\mathcal{G}}(y)\rangle$ by bounding its higher order moments. Fix a hash function $h \in_u \mathcal{H}$ and for $j \in \{1, \ldots, m\}$, let

$$Z_j = \langle v_{|h^{-1}(j)},\ \mathcal{G}'(z_j)\rangle.$$

Note that the random variables Z_j are independent of one another, and

$$\langle v,\ \bar{\mathcal{G}}(y)\rangle = \sum_{j=1}^{m} Z_j.$$

We use Lemma 10 to bound $\left\|v_{|h^{-1}(j)}\right\|_2$. We defer the proof of the following technical lemma.

Lemma 18. *With probability $1 - 2m\sqrt{\gamma}$, for all $j \in [m]$ we have*

$$\left\|v_{|h^{-1}(j)}\right\|_2^2 \leq \frac{2}{m} \tag{8}$$

We condition on the hash function h satisfying Equation (8), and call this event A.

Recall that $Z_j = \langle v_{|h^{-1}(j)},\ \mathcal{G}'(z_j)\rangle$. By Lemma 12, with probability $1 - \gamma$ over z_j, we have the bound

$$|Z_j| \leq d(n,\gamma)\left\|v_{|h^{-1}(j)}\right\|_2 \leq \frac{\sqrt{2}d(n,\gamma)}{\sqrt{m}} \leq \frac{1}{2\sqrt{\log(1/\gamma)}} := M \tag{9}$$

where the last inequality is by the choice of m. By the union bound, Equation (9) holds with probability at least $1 - m\gamma$ over $z_1, \ldots, z_m$, for all $j \in [m]$. We further condition on the event $|Z_j| \leq M$ for all $j \in [m]$ which we denote by B.

Conditioning on a high probability event preserves the small-bias property of $G(z_j)$'s up to a small additive error. In particular, conditioned on the event B, $\mathcal{G}'(z_j)$ is $(\varepsilon + m\gamma)$-biased. Since $\left\|v_{|h^{-1}(j)}\right\|_1 \leq \sqrt{n}$ we have

$$\begin{aligned}\sum_{j=1}^{m} \mathbb{E}[Z_j^2|B] &\leq \sum_{j=1}^{m}(\left\|v_{|h^{-1}(j)}\right\|_2^2 + n(\varepsilon + m\gamma)) \\ &\leq 1 + nm(\varepsilon + m\gamma) \\ &\leq 2.\end{aligned}$$

Further, since $\mathcal{G}'(z_j)$ is symmetric, it continues to be symmetric after we condition on B (which is a symmetric event in $\mathcal{G}'(z_j)$).

Since $t/2 \leq \sqrt{\log(1/\gamma)}$, and $M = 1/2\sqrt{\log(1/\gamma)}$ we have $Mt \leq 1/2$. We now apply Bernstein's inequality [Fel71] to the random variables $\{Z_j|B\}_{j=1}^m$ which are mean zero and are bounded by M to get

$$\begin{aligned}\Pr\left[\left|\sum_{j=1}^m Z_j|B\right| > t/2\right] &\leq 2\exp\left(-\frac{t^2}{4(\sum_j \|Z_j|B\|_2^2 + Mt/3)}\right)\\ &\leq 2\exp\left(-\frac{t^2}{28/3}\right)\\ &\leq 2\exp\left(-\frac{t^2}{16}\right).\end{aligned}$$

Combining the above arguments, we get that

$$\Pr\left[\left|\langle v,\ \mathcal{G}'(y)\rangle\right| > t/2\right] \leq 2\exp(-t^2/16) + 2m\sqrt{\gamma}. \tag{10}$$

The claim now follows from combining Equations (10) and (7) □

Proof of Lemma 18. Note that

$$\|v \star v\|_1 \leq 1,\ \|v \star v\|_2 = \|v\|_4^2 \leq \beta.$$

Therefore, by setting

$$p = \frac{\log(1/\gamma)}{\log(1/\beta)},\ t = D_6\sqrt{\beta p}$$

in Lemma 10 we get

$$\Pr_{h\in\mathcal{H}}\left[\left\|v_{|h^{-1}(j)}\right\|_2^2 \geq \frac{1}{m} + D_6\sqrt{\beta\frac{\log(1/\gamma)}{\log(1/\beta)}}\right] \leq \frac{D_6^p p^p \beta^p + \gamma}{(D_6\sqrt{\beta p})^p} \leq 2\sqrt{\gamma}.$$

By a union bound, with probability at least $1 - 2m\sqrt{\gamma}$ over h, for all $j \in [m]$,

$$\begin{aligned}\left\|v_{|h^{-1}(j)}\right\|_2^2 &\leq \frac{1}{m} + \sqrt{\beta\frac{\log(1/\gamma)}{\log(1/\beta)}}\\ &\leq \frac{2}{m}\end{aligned}$$

where the last inequality follows from our choice of β in Equation (6). □

3.3 Putting things together

We are now ready to prove Theorem 11.

Proof of Theorem 11. Let $\bar{\mathcal{G}} : \left(\{0,1\}^{r'}\right)^m \times \mathcal{H} \to \{\pm 1\}^n$ be the generator as in Lemma 17. Let δ be the final additive error desired. Set

$$\gamma = \frac{\delta^2}{(D_5 \log(n/\delta))^{D_6 \log\log(n)}} \leq \frac{\delta^2}{50m^2}$$
$$\varepsilon = \left(\frac{\delta}{n}\right)^{D_7 \log\log(n/\delta)^3} \leq \frac{\delta}{20\gamma^{4\log(m)}}.$$

It can be verified that with these parameter seedings, the error probability in Lemma 17 is at most $4e^{-t^2/16} + \delta$, and the seed-length of $\bar{\mathcal{G}}$ is $\log(n/\gamma) + mr'$, where $r' = O(\log(n/\delta) \cdot (\log\log(n/\delta))^3)$.

Observe that once we fix the has function h, the inner product $\langle w,\ \mathcal{G}'(z_1, \ldots, z_m, h)\rangle$ can be computed by a (S, r', m)-ROBP where $S = O(\log n)$ which reads one $z_1, z_2, \ldots, z_m$ in order. The reason is that we can round each weight w_i up to a multiple of $1/n^2$. This can only increase $|\langle w,\ z\rangle|$ by $1/n$ for any $z \in \{\pm 1\}^n$. This also ensures that $\langle w,\ z\rangle$ lies in the interval $[-\sqrt{n}, \sqrt{n}]$ and that it is a multiple of $1/n^2$ and thus can be computed with $O(\log n)$-bits of precision.

Now, let $\mathcal{G}^{INW} : \{0,1\}^{r_s} \to \left(\{0,1\}^{r'}\right)^m$ be a generator fooling $O(S, r', m)$-ROBP with error δ. By Theorem 5, there exist such generators with seed-length

$$r_s = O(r' + (\log n)(\log m) + \log(m/\delta) \cdot (\log m)) = O((\log(n/\delta))(\log\log(n/\delta))^3).$$

Now, if we define our final generator $\mathcal{G}^f : \mathcal{H} \times \{0,1\}^{r_s} \to \{\pm 1\}^n$ by

$$\mathcal{G}^f(h, z) = \mathcal{G}'(\mathcal{G}^{INW}(z), h).$$

From the above arguments it follows that the output of $\mathcal{G}^f$ only has an additional δ error compared to $\mathcal{G}'$. The theorem now follows from the above bound on seed-length. □

4 A PRG for signed majorities

In this section we construct generator to fool signed majorities to polynomial error with seedlength $\tilde{O}(\log n)$ proving Theorem 2. As the generator and its analysis is quite technical, we first give a high-level description at the risk of repeating parts of Section 1.3.

Proof overview

For simplicity, in this discussion let us fix a test vector $v \in \{-1, 0, 1\}^n$ and error $\epsilon = 1/\mathrm{poly}(n)$. We start by noting that it suffices to design a PRG $G : \{0,1\}^r \to \{\pm 1\}^n$ such that $d_{TV}(\langle v, G(y)\rangle, \langle v, X\rangle) \ll 1/\mathrm{poly}(n)$ where $y \in_u \{0,1\}^r$ and $X \in_u \{\pm 1\}^n$. In the following let $X \in_u \{\pm 1\}^n$ and $Y \sim G(y)$, where $y \in_u \{0,1\}^r$ be the output of the desired generator.

The starting point of our analysis and construction is to note that showing closeness in statistical distance for discrete random variables is equivalent to showing that the *Fourier transforms* of the random variables are close. This will allow us to use various analytic tools. Concretely, we shall use the following elementary fact about the discrete Fourier transform.

Claim 2. *Let Z_1, Z_2 be two discrete random variables with support sizes at most B. Then,*

$$d_{TV}(Z_1, Z_2) \leq \sqrt{2B} \cdot \max_{\alpha \in \mathbb{R}} \left| \mathbb{E}[\exp(2\pi i \alpha Z_1)] - \mathbb{E}[\exp(2\pi i \alpha Z_2)] \right|.$$

Proof. Note that the distribution $Z_1 - Z_2$ is supported on at most $2B$ points. Therefore,

$$d_{TV}(Z_1, Z_2) = \|Z_1 - Z_2\|_1 \leq \sqrt{2B} \|Z_1 - Z_2\|_2.$$

On the other hand, the Plancherel identity implies that

$$\|Z_1 - Z_2\|_2 \leq \max_{\alpha \in \mathbb{R}} \left| \mathbb{E}[\exp(2\pi i \alpha Z_1)] - \mathbb{E}[\exp(2\pi i \alpha Z_2)] \right|.$$

This completes the proof. □

Henceforth, we will focus on designing a generator so as to fool the test function $\exp(2\pi i \alpha \langle v, x\rangle) \equiv \phi_{v,\alpha}(x)$. To do so, we will consider two cases based on how large $\alpha \in [0,1]$ is. The two cases we consider capture the shift in the behaviour of $\mathbb{E}[\exp(2\pi\alpha(v \cdot X))]$ - the "α-th Fourier coefficient". We can combine the generators for the two cases easily at the end.

Large α: $\alpha \gg (\log n)^{O(1)}/\sqrt{\|v\|_0}$

Roughly speaking, the reason for considering this threshold is that all values of α greater than this value yield similar Fourier coefficients: $|\mathbb{E}[\phi_{v,\alpha}(X)]| \ll 1/\mathrm{poly}(n)$ for α in this range. Thus, it suffices to ensure that $\mathbb{E}[\phi_{v,\alpha}(Y)]$ is small. We achieve this by exhibiting a way to "amplify" the error, i.e., go from fooling $\phi_{v,\alpha}$ with constant error to fooling them with polynomially small error at the expense of a $O(\log \log n)$ factor in seed-length. We then instantiate this amplification procedure with the generator of Gopalan, Meka, Reingold, Zuckerman [GMRZ13] which requires seed-length $O(\log n)$ to fool such test functions $(\phi_{v,\alpha}(\ ))$ with constant error. We leave the details of the amplification procedure to the corresponding section.

Small α: $\alpha \ll (\log n)^{O(1)}/\sqrt{\|v\|_0}$

This is the harder of the two cases and the core of our construction and analysis. The generator we use is essentially the same as the one based on iterative dimension reduction used in derandomizng the Chernoff bound. The main difference will be that instead of using small-bias spaces in each dimension reduction step we use k-wise independent spaces for suitable k. However, the analysis is quite different and requires several new analytic tools.

We next formally describe our generator for handling this case. Let $n, \delta > 0$. Let C be a sufficiently large constant. We define a generator as follows. Let $n = n_1 > n_2 > \ldots > n_t$ so that $n_{i+1} = n_i^{1/2} + O(1)$ and $\log^{2C}(n/\delta) \geq n_t \geq \log^C(n/\delta)$. Note that this implies that $t = O(\log\log(n))$. For $1 \leq i < t$, let $\mathcal{H}_i = \{h : [n_i] \to [n_{i+1}]\}$ be a family of $\frac{C\log(n/\delta)}{\log(n_i)}$-wise independent hash functions. Let $h_i \in_u \mathcal{H}_i$. Let Z_i be a random element of $\{\pm 1\}^{n_i}$ chosen from a distribution that is both $(\delta/n)^C$-biased and and $\frac{C\log(n/\delta)}{\log(n_i)}$-wise independent. Finally, let Z be a random variable in $\{\pm 1\}^{n_t}$ be chosen to fool weight at most n halfspaces to variational distance δ/n as described in Theorem 4. We define our random variable $Y \in \{\pm 1\}^n$ to be

$$Y = ZA(h_{t-1})D(Z_{t-1})A(h_{t-2})D(Z_{t-2})\cdots A(h_1)D(Z_1). \tag{11}$$

Informally, this generator begins with the string Z_1, then uses h_1 to divide the coordinates into n_2 bins and then for each bin multiplies the elements in this bin by a random sign, these n_2 signs being chosen recursively by a similar generator, until at the final level they are picked using the generator from Theorem 4 instead.

It is easy to see from Theorem 4 and Fact 3 that the random variable Y can be produced from a random seed of length $s = O(\log(n/\delta)\log\log(n/\delta))$. We also claim that it fools $\phi_{v,\alpha}$ for $|\alpha| \leq \log^3(1/\delta)/\|v\|_2$. This in turn implies our claimed pseudorandomness for halfspaces in lieu of Claim 2.

To analyze the generator we shall use a hybrid argument to exploit the recursive nature of the generator. To this end, for $1 \leq i < t$, let $X_i \in_u \{\pm 1\}^{n_i}$ and define

$$Y_i := X_i A(h_{i-1})D(Z_{i-1})\cdots A(h_1)D(Z_1) \tag{12}$$

(note that $Y_1 = X_1$) and let $Y_t = Y$.

The crux of the analysis is then in showing the following claim analyzing a single dimension reduction step: for $1 \leq i \leq t$ and $\alpha \leq \log^3(1/\delta)/\|v\|_2$,

$$|\mathbb{E}[\phi_{v,\alpha}(Y_i)] - \mathbb{E}[\phi_{v,\alpha}(Y_{i+1})]| \leq \delta/n.$$

If we let $v_0 = v$ and $v_i = A(h_{i-1})D(Z_{i-1}\cdots A(h_1)D(Z_1)v$, then the above claim amounts to bounding

$$|\mathbb{E}[\phi_{v_i,\alpha}(X_i)] - \mathbb{E}[\phi_{v_i,\alpha}(X_{i-1}A(h_i)D(Z_i))]|\,. \tag{13}$$

Thus, intuitively, we need to argue that a single step of dimension reduction (i.e., applying $A(h_i)D(Z_i)$) does not cause too much error. Ideally, we would have liked to make such a claim for all test functions of the form $\phi_{w,\alpha}$; this turns out to be false. What remains true however is that a single dimension reduction step *fools* test functions of the form $\phi_{w,\alpha}$ when the test vector $w \in \mathbb{R}^{n_i}$ is *sufficiently well-spread out* (as measured by the ℓ_2, ℓ_4-norms of w) and α is not too large. In particular, in the most technically intensive part of our argument we bound the error from the above step as a function of the ℓ_2, ℓ_4 norms of the vector v_i. We then argue separately that the ℓ_2, ℓ_4 norms of the test vector v are close to their true values under the above transformations.

In order to analyze expectations as in Equation 13, it is critical to note that X_{i-1} is uniformly distributed. This implies (for fixed h_i) that the given expectation over X_{i-1} is a product of cosines of linear functions of Z_i. We take advantage of the fact that cosine is a smooth function of its input, allowing us to approximate this product by a Taylor polynomial. If α is sufficiently small, the higher order terms will be small enough to ignore, and therefore the limited independence of Z_i will be sufficient to guarantee the desired approximation.

4.1 Generator for large α

We now develop a generator that works when α is large, in particular, we prove:

Proposition 19. *There exists an explicit generator $G^b : \{0,1\}^r \to \{\pm 1\}^n$ with seed-length $r = O((\log(n/\epsilon))(\log\log(n)))$ such that the following holds. For all $v \in \{-1, 0, 1\}^n$, $\alpha \in (-1/4, 1/4)$ with $\alpha \geq \log^3(1/\epsilon)/\left\|v\right\|_2$,*

$$\left| \mathop{\mathbb{E}}_{y \in_u \{0,1\}^r} \left[\phi_{v,\alpha}(G^b(y)) \right] - \mathop{\mathbb{E}}_{X \in_u \{\pm 1\}^n} \left[\phi_{v,\alpha}(X) \right] \right| \leq \epsilon.$$

4.1.1 Spreading hashes

In order to prove Proposition 19 we will need to study a certain property of hash families.

Definition 3. *A family of hash functions $\mathcal{H} = \{h : [n] \to [m]\}$ is said to be (k, ℓ, ϵ)-spreading if the following holds: for every $I \subseteq [n]$ with $|I| \geq k$, and $h \in_u \mathcal{H}$ with probability at least $1 - \epsilon$, then for all $j \in [m]$, $|h^{-1}(j) \cap I| \leq |I|/\ell$.*

The above definition quantifies the intuition that when a sufficiently large (so that standard tail bounds apply) collection of items $I \subseteq [n]$ is hashed into m bins, the max-load is not much more than the average load of $|I|/m$. It will be important for us to be able to construct such families explicitly.

Lemma 20. *For all $\epsilon \geq 0$, there exists an explicit hash family $\mathcal{H} = \{h : [n] \to [m]\}$ where $m = O(\log^5(1/\epsilon))$ which is $((\log^5(1/\epsilon)), \log(1/\epsilon), \epsilon)$-spreading and $h \in_u \mathcal{H}$ can be sampled with $O(\log(n/\epsilon))$ bits.*

Proof. Let $m = \Theta(\log^5(1/\epsilon))$ and let $\mathcal{H} = \{h : [n] \to [m]\}$ be a δ-biased family for $\delta = \exp(-C(\log(1/\epsilon)))$ for C a sufficiently large constant. We argue that $\mathcal{H}$ satisfies the conditions of the lemma by standard moment bounds.

Let $p = 2\log(1/\epsilon)/\log\log(1/\epsilon))$. Let $|I| > \log^5(1/\epsilon)$ and let $v \in \{0,1\}^n$ be the indicator vector of the set I. Note that if some h has $|h^{-1}(j) \cap I| > |I|/\log(1/\epsilon)$ for some j, then $h(v) \geq |I|^2/\log^2(1/\epsilon)$ (recall the definition of $h(v)$ from Equation 1). Therefore, by Lemma 9 and Markov's inequality, the probability that this happens is at most

$$\begin{aligned}\frac{\mathbb{E}[h(v)^p]\log^{2p}(1/\epsilon)}{|I|^{2p}} &\leq O\left(\frac{p^2\log^2(1/\epsilon)}{m} + p^2\log^2(1/\epsilon)|I|^{-1}\right)^p + m^p\log^{2p}(1/\epsilon)\delta \\ &= O(\log(1/\epsilon))^{-p} + O(\log(1/\epsilon))^{5p}\delta \\ &\leq \epsilon.\end{aligned}$$

□

4.1.2 The PRG

We begin with a simpler version of our generator which has the desired pseudorandomness property but has too large a seed. We will then improve the seed-length using PRGs for small-space machines.

Let $\mathcal{H} = \{h : [n] \to [m]\}$ be a $(k, C\log(1/\epsilon), \epsilon)$-spreading family for parameters k, C, ϵ to be chosen later. Let $G^{CS} : \{0,1\}^r \to \{\pm 1\}^n$ be a generator as in Theorem 4 with error $1/4$. Now, define the generator $G^b : \mathcal{H} \times (\{0,1\}^r)^m \to \{\pm 1\}^n$ as follows: for $i \in [n]$,

$$G^b(h, z_1, z_2, \ldots, z_m)_i = G^{CS}(z_{h(i)})_i. \tag{14}$$

We claim that the above generator fools tests of the form $\phi_{v,\alpha}(\ )$ for $\|v\|_0 \geq k$ and $\alpha \gg \sqrt{m}/\|v\|_2$.

Lemma 21. *Let C be a sufficiently large constant. Let $\mathcal{H} = \{h : [n] \to [m]\}$ for some $m \geq \log(1/\epsilon)$ be a $(k, \ell, \epsilon/4)$-spreading family with $\ell = C\log(1/\epsilon)$. Let G^{CS} be a generator as in Theorem 4 with error $1/4$. Let $Y \in \{\pm 1\}^n$ be the output of the generator G^b as defined in Equation (14) on a uniformly random seed and $X \in_u \{\pm 1\}^n$. Then, for all $v \in \{-1, 0, 1\}^n$ with $\|v\|_0 \geq k$, and $C\sqrt{m}/\|v\|_2 \leq \alpha \leq 1/4$,*

$$|\mathbb{E}[\phi_{v,\alpha}(Y)] - \mathbb{E}[\phi_{v,\alpha}(X)]| \leq \epsilon.$$

Proof. Fix the test vector $v \in \{-1, 0, 1\}^n$. Let $I = Supp(v)$ and let $|I| = K \geq k$. Let $Y = G^b(h, z_1, z_2, \ldots, z_m)$ and for $j \in [m]$, let $Y^j = G^{CS}(z_j)$ and let $X^j \in_u \{\pm 1\}^n$ be independent uniformly random strings. Suppose that the hash function $h \in_u \mathcal{H}$ is such that the condition of $(k, \ell, \epsilon/4)$-spreading holds for I. This assumption only incurs an additive $\epsilon/2$ in the error.

First note that,

$$\mathbb{E}[\phi_{v,\alpha}(X)] = (\cos 2\pi\alpha)^K \leq \exp(-\Omega(\alpha^2 K)) \leq \exp(-Cm) \leq \epsilon/4.$$

Thus, we need only show that

$$\mathbb{E}[\phi_{v,\alpha}(Y)] \leq \epsilon/4.$$

Now, for $j \in [m]$ let $v^j = v_{h^{-1}(j)}$ and $K_j = |I \cap h^{-1}(j)|$. Observe that by definition, $d_{TV}(v^j \cdot Y^j, v^j \cdot X^j) \leq 1/4$. Therefore,

$$\left|\mathbb{E}[\phi_{v^j,\alpha}(Y^j)] - \mathbb{E}[\phi_{v^j,\alpha}(X^j)]\right| \leq 1/2.$$

Further,

$$\mathbb{E}[\phi_{v^j,\alpha}(X^j)] = (\cos 2\pi\alpha)^{K_j} = \exp(-\Omega(\alpha^2 K_j)).$$

Combining the above two equations, we get

$$|\mathbb{E}[\phi_{v,\alpha}(Y)]| = \left|\prod_{i=1}^{m} \mathbb{E}[\phi_{v^j,\alpha}(Y^j)]\right| \leq \prod_{i=1}^{m} \min\left(\left(\frac{1}{2} + \exp(-\Omega(\alpha^2 K_j))\right), 1\right).$$

Now, because h has the well-spreading property, $K_j = |h^{-1}(j) \cap I| \leq |I|/\ell$ for all $j \in [m]$. On the other hand, $\sum_j K_j = K$. Since the sum of the K_j which are at most $K/(2m)$ totals at most $K/2$ and since none of the other K_j are too large, there must be at least $\ell/2$ values of j so that $K_j \geq K/(2m)$. For these j we have that

$$\frac{1}{2} + \exp(-\Omega(\alpha^2 K_j)) \leq \frac{1}{2} + \exp(-\Omega((C^2 m/K)(K/2m))) = \frac{1}{2} + \exp(-\Omega(C)) \leq \frac{3}{4}$$

for C sufficiently large. Thus, for C sufficiently large

$$|\mathbb{E}[\phi_{v,\alpha}(Y)]| \leq \left(\frac{3}{4}\right)^{\ell/2} \leq \epsilon/4.$$

This completes the proof.

□

We are now ready to prove Proposition 19.

Proof of Proposition 19. Let C be a sufficiently large constant, $m = C\log^5(1/\epsilon)$, let $\mathcal{H} = \{h : [n] \to [m]\}$ be a $(k, \ell, \epsilon/4)$-spreading family with $k \leq C\log^5(1/\epsilon)$, and $\ell = C\log(1/\epsilon)$ as given in Lemma 20. Note that if $1/4 \geq \alpha \geq \log^3(1/\epsilon)/\|v\|_2$ for some α, it must be the case that $\|v\|_0 \geq \log^6(1/\epsilon) \geq k$. Therefore, Lemma 21 provides us with a generator, G^b, so that for any such α that if Y is an output of G^b and X a uniform random element of $\{\pm 1\}^n$ and if $\|v\|_0 \geq k$, then

$$|\mathbb{E}[\phi_{v,\alpha}(Y)] - \mathbb{E}[\phi_{v,\alpha}(X)]| \leq \epsilon/2.$$

Unfortunately, the seed-length of G^b is $\log(\mathcal{H}) + O(\log n) \cdot m$. We improve this using the PRGs for ROBPs of Theorem 5. It is easy to see that for a fixed hash function h and test vector v, the computation of $\langle v, G^b(h, z_1, \ldots, z_m)\rangle$ can be done by a (S, D, m)-ROBP where $S = O(\log n)$ and $D = O(\log n)$. Thus, we can further derandomize the choice of $z_1, \ldots, z_m$ using the PRG from Theorem 5. Formally, let $G^{INW} : \{0,1\}^r \to \left(\{0,1\}^D\right)^m$ be a generator fooling (S, D, m)-ROBPs as in Theorem 5 with error $\epsilon/4$ and define

$$G^f(h, z) = G^b(h, G^{INW}(z)).$$

Then, from the above arguments it follows that G^f fools $\phi_{v,\alpha}$ with error at most ϵ and has seed-length $O(\log(n/\epsilon) \cdot (\log\log(n/\epsilon)))$ proving the claim. □

4.2 Generator for small α

We next argue that the generator defined in Equation 11 fools Fourier coefficients $\phi_{v,\alpha}$ for sufficiently small α. The main claim of this section is the following.

Proposition 22. *Let $v \in \{-1, 0, 1\}^n$ and $\alpha \in \mathbb{R}$ with $|\alpha| \leq \log^3(1/\delta)/\|v\|_2$. Let C be a sufficiently large constant and let $\delta > 0$. Let Y be as defined by Equation (11) and let $X \in_u \{\pm 1\}^n$. Then*

$$|\mathbb{E}[\phi_{v,\alpha}(Y)] - \mathbb{E}[\phi_{v,\alpha}(X)]| \leq \delta.$$

As described in the overview, we will prove the claim by a hybrid argument. For ease of notation, we repeat some notation from the overview section. For $1 \leq i < t$, letting X_i be a uniform random element of $\{\pm 1\}^{n_i}$ we define

$$Y_i := X_i A(h_{i-1}) D(Z_{i-1}) \cdots A(h_1) D(Z_1) \tag{15}$$

(note that $Y_1 = X_1$) and let $Y_t = Y$. Our Proposition will follow from the following Lemma.

Lemma 23. *With Y_i defined as above for C sufficiently large and $v \in \{-1, 0, 1\}^n$ and $\alpha \in \mathbb{R}$ with $|\alpha| \leq \log^3(1/\delta)/\|v\|_2$, then for $t > i \geq 1$*

$$|\mathbb{E}[\phi_{v,\alpha}(Y_{i+1})] - \mathbb{E}[\phi_{v,\alpha}(Y_i)]| \leq \delta/n.$$

The proof of the Lemma 23 will be further split into two main cases based upon whether or not the vector v is sparse relative to n_i. Intuitively, the case of sparse v is easier as hashing takes care of most issues here.

4.2.1 Analysis for sparse vectors

We begin with the case where v is sparse.

Lemma 24. *With $Y_i, C, n, v, \alpha, \delta$ as in Lemma 23 with $i < t$, if $\|v\|_0^3 < n_{i+1}$ then*

$$|\mathbb{E}[\phi_{v,\alpha}(Y_{i+1})] - \mathbb{E}[\phi_{v,\alpha}(Y_i)]| \leq \delta/n.$$

Proof. We claim that this holds even after fixing the values of h_j, Z_j for all $j < i$. In particular, if we let

$$w = vD(X_1)A(h_1)^T \cdots D(h_{i-1})A(h_{i-1})^T$$

then we need to show that

$$|\mathbb{E}[\phi_{w,\alpha}(X_{i+1}A(h_i)D(Z_i))] - \mathbb{E}[\phi_{w,\alpha}(X_i)]| \leq \delta/n.$$

We will show the stronger claim that

$$d_{TV}(w \cdot X_{i+1}A(h_i)D(Z_i), w \cdot X_i) \leq \delta/(2n).$$

Intuitively, this will hold because v (and hence w) is sparse. This means that with high probability h_i will cause few collisions within the support of w. If this is the case, then Z_i will nearly randomize the relative signs of elements mapped to the same bin and X_i will randomize the signs between bins. To show that we have few collisions, we will need the following lemma:

Lemma 25. *Let n and m be positive integers, $\epsilon > 0$ and C a sufficiently large constant. Let $\mathcal{H} = \{h : [n] \to [m]\}$ be a k-wise independent family of hash functions for $k = \frac{C\log(m/\epsilon)}{\log(m)}$. Let $I \subset [n]$ be such that $|I|^3 \leq m$. Then for $h \in_u \mathcal{H}$, with probability at least $1 - \epsilon$ we have that*

$$|I| - |h(I)| \leq k.$$

Proof. Note that if $|I| - |h(I)| > k$ then at least k elements of I were sent to the same location as some other element of I. This implies that there must be at least $k/3$ disjoint pairs of elements $x_i, y_i \in I$ so that $h(x_i) = h(y_i)$ (for each element $j \in [m]$ so that $|h^{-1}(j)| = \ell > 1$ we can find at least $\ell/3$ pairs). Thus, it suffices to show that the expected number of collections of distinct elements $x_1, y_1, x_2, y_2, \ldots, x_{k/3}, y_{k/3} \in I$ so that $h(x_i) = h(y_i)$ for each i is less than ϵ. On the other hand, the number of sequences $x_i, y_i \in I$ is at most $|I|^{2k/3}$ and the probability that any given sequence has the desired property is $m^{-k/3}$ by k-wise independence of h. Thus the expected number of such sets of pairs is at most

$$|I|^{2k/3}m^{-k/3} \leq m^{2k/9}m^{-k/3} = m^{-k/9} \leq \epsilon.$$

This completes the proof. □

Applying this lemma to $I = \text{supp}(w)$, we find that except with probability $\delta/(4n)$ we have that at most $\log(n/\delta)$ elements of I collide with any other element of I under h_i. Let J be the set of such coordinates. It is clear that the distribution of $w \cdot (X_{i+1}A(h_i)D(Z_i))$ as we vary X_{i+1} depends only on h_i and the signs of the Z_i on the coordinates of J. On the other hand, it is easy to see that the restriction of Z_i to these coordinates is within $2^{|J|}(\delta/n)^C < \delta/(4n)$ of uniform. Thus,

$$\delta/(2n) \geq d_{TV}(w\cdot(X_{i+1}A(h_i)D(Z_i)), w\cdot(X_{i+1}A(h_i)D(X_i))) = d_{TV}(w\cdot X_{i+1}A(h_i)D(Z_i), w\cdot X_i).$$

Where the equality above is because $(X_{i+1}A(h_i)D(X_i))$ and X_i are identically distributed. This completes the proof. □

4.2.2 Analysis for dense vectors

For relatively dense vectors v, we will need a different, more analytic approach. The following crucial lemma analyzes the effect of a single dimension reduction step and bounds the error in terms of the norms of the test vector v. We will then apply the lemma iteratively.

Lemma 26. *Let $\delta > 0$, $n, m \geq 1$, and $p \geq 2$ and even integer. Let $\mathcal{D}$ be a $2p$-wise independent distribution over $\{\pm 1\}^n$ and $\mathcal{H} = \{h : [n] \to [m]\}$ be a $2p$-wise independent hash family. Then, for all $v \in \mathbb{R}^n$, $X \in_u \{\pm 1\}^n$,$Y \sim \mathcal{D}$, $h \in_u \mathcal{H}$ and $Z \in_u \{\pm 1\}^m$,*

$$|\mathbb{E}[\phi_{v,\alpha}(Z \cdot A(h) \cdot D(Y))] - \mathbb{E}[\phi_{v,\alpha}(Z \cdot A(h) \cdot D(X))]| < O(p)^{2p} \left(\frac{\alpha^4 \|v\|_2^4}{m} + \alpha^4 \|v\|_4^4 \right)^{p/8}. \tag{16}$$

To prove the lemma we shall exploit the independence of the z_i's in Equation (18) to reduce the problem to that of analyzing a product of cosines as in the following lemma. The lemma gives a low-degree (multivariate) polynomial approximation for a product of cosines.

Lemma 27. *For all $\alpha \in (0, 1/4)$ and even integer p, there exists a polynomial $P : \mathbb{R}^m \to \mathbb{R}$ of degree at most p such that for all $S_1, \ldots, S_m, T \in \mathbb{R}$,*

$$\prod_{j=1}^{m} \cos(2\pi\alpha S_j) = \exp(-2\pi^2\alpha^2 T) \cdot \left(\sum_{t=0}^{p/2-1} \frac{\left(-2\pi^2\alpha^2 \left(\sum_{i=1}^m S_i^2 - T\right)\right)^t}{t!} P(S_i) \right) \tag{17}$$

$$+ O(1)^p \left(\left(\alpha^2 \left(\sum_{i=1}^m S_i^2 - T \right) \right)^{p/2} + \left(\alpha^2 \left(\sum_{i=1}^m S_i^2 - T \right) \right)^{p} + \left(\sum_{i=1}^m (\alpha S_i)^4 \right)^{p/8} + \left(\sum_{i=1}^m (\alpha S_i)^4 \right)^{p/2} \right).$$

Proof of Lemma 26. Let $Y^s = ZA(h)D(Y)$. We first fix a hash function $h \in \mathcal{H}$ and then bound the error as a function of the hash function. We then average the error bound for a uniformly random hash function from $\mathcal{H}$ using Lemma 9.

For $j \in [m]$, let random variable $S_j = \sum_{i:h(i)=j} v_i Y_i$. Note that $\langle v,\ Y^s \rangle = \sum_{j=1}^m z_j S_j$. Therefore, as $z \in_u \{\pm 1\}^m$,

$$\mathop{\mathbb{E}}_{z}[\phi_{v,\alpha}(Y^s)] = \prod_{j=1}^{m} \cos(2\pi\alpha S_j).$$

Let $T = \sum_j \mathbb{E}[S_j^2] = \|v\|_2^2$. Let $Q(X) \equiv R(S_1, \ldots, S_m)$ denote the degree $2p$ polynomial corresponding in the first term of Equation (17), and let $E(X)$ be the *error* term corresponding to the second term. Then, from the above calculations,

$$\mathop{\mathbb{E}}_{z}[\phi_{v,\alpha}(X^s)] = Q(X) + E(X).$$

Observe that

$$Q_2(X) := \alpha^2(\sum_{j=1}^{m} S_j^2 - T) = \alpha^2 \sum_{j=1}^{m} \sum_{i \neq i' \in h^{-1}(j)} v_i v_{i'} X_i X_{i'}, \tag{18}$$

is a degree two polynomial in X with, $\|Q_2\|_2^2 \leq \alpha^4 h(v)$ (recall Equation (1)).

By hypercontractivity - Lemma 6, for all even $r \leq p$,

$$\mathbb{E}[Q_2(X_1, \ldots, X_n)^r] \leq O(r)^r \left(\alpha^4 h(v)\right)^{r/2}.$$

A similar calculation for the polynomial $Q_4(X_1, \ldots, X_n) := \alpha^4(\sum_j S_j^4)$ shows that for all even $r \leq p/2$,

$$\mathbb{E}[Q_4(X_1, \ldots, X_n)^r] \leq O(r)^{2r} \left(\alpha^4 h(v)\right)^{r}.$$

By $2p$-wise independence, the above bounds also hold for $\mathbb{E}[Q_2(Y)^r], \mathbb{E}[Q_4(Y)^r]$.

Now, let $X^s = ZA(h)D(X)$, where $X \in_u \{\pm 1\}^n$. Then, clearly $X^s \in_u \{\pm 1\}^n$. Combining the above expressions and noting that they also work for $X \in_u \{\pm 1\}^n$, we get

$$\begin{aligned}
\left|\mathop{\mathbb{E}}_{X}\mathop{\mathbb{E}}_{z}[\phi_{v,\alpha}(X^s)] - \mathop{\mathbb{E}}_{Y}\mathop{\mathbb{E}}_{z}[\phi_{v,\alpha}(Y^s)]\right| &\leq |\mathbb{E}\left[Q(X) - Q(Y)\right]| + \mathbb{E}[|E(X)|] + \mathbb{E}[|E(Y)|] \\
&\leq 0 + O(p)^p \left(\alpha^4 h(v)\right)^{p/4} + O(p)^p \left(\alpha^4 h(v)\right)^{p/8} \\
&\leq O(p)^p \left(\alpha^4 h(v)\right)^{p/8}.
\end{aligned}$$

By taking expectation with respect to $h \in_u \mathcal{H}$ and applying Lemma 9, we get

$$|\mathbb{E}[\phi_{v,\alpha}(X^s)] - \mathbb{E}[\phi_{v,\alpha}(Y^s)]| \leq O(p)^{2p} \left(\frac{\alpha^4 \|v\|_2^4}{m} + \alpha^4 \|v\|_4^4\right)^{p/8},$$

proving the lemma. □

We defer the proof of Lemma 27 to Section 4.2.5 and continue with the analysis of our generator. We do so by applying Lemma 26 iteratively to the vectors

$$v_i := vD(Z_1)A(h_1)^T \cdots D(Z_{i-1})A(h_{i-1})^T.$$

In order for it to be useful, we need to have good bounds on the low order moments of the v_i. We deal with these issues in the next section.

4.2.3 Controlling moments

In particular we will need the following Lemma:

Lemma 28. *Let $v \in \{-1, 0, 1\}^n$ with $\|v\|_0 \geq \log^{C/4}(n/\delta)$. Let Z_i, h_i, v_i be defined as above. For any $1 \leq i \leq t$ we have with probability at least $1 - \delta/(4n)$ that*

$$\|v_i\|_2 \leq 2^i \|v\|_2 \qquad \text{and} \qquad \|v_i\|_4 \leq \frac{\|v\|_2}{\min(\|v\|_2^{1/3}, n_i^{1/20})}.$$

In order to prove this we will first need some controls over how the procedure used to obtain v_{i+1} from v_i affects these norms. In particular, we show:

Lemma 29. *Let $p \geq 2$ be an even integer. Let $\mathcal{H} = \{h : [n] \to [m]\}$ be a $4p$-wise independent hash family and $\mathcal{D}$ be a $4p$-wise independent distribution over $\{\pm 1\}^n$. Then, for $h \in_u \mathcal{H}$, $x \sim \mathcal{D}$ and a vector $v \in \mathbb{R}^n$,*

$$\mathbb{E}\left[\left(\|v\|_2^2 - \left\|vD(x)A(h)^T\right\|_2^2\right)^p\right] \leq O(p)^{2p}\left(\frac{\|v\|_2^4}{m}\right)^{p/2} + O(p)^{2p}\|v\|_4^{2p}.$$

Similarly,

$$\mathbb{E}[\left\|vD(x)A(h)^T\right\|_4^{4p}] \leq O(p)^{4p}\left(\frac{\|v\|_2^4}{m}\right)^{p} + O(p)^{4p}\|v\|_4^{4p}.$$

Proof. Note that in either case the independence is sufficient that the expectations would be the same if x and h were chosen uniformly at random from $\{\pm 1\}^n$ and $[m]^{[n]}$, respectively.

Applying Lemma 6 to the polynomial $P_h(x) = \|vD(x)A(h)^T\|_2^2 - \|v\|_2^2$, we find that for fixed h

$$\mathop{\mathbb{E}}_x[P_h(x)^p] \leq O(p)^p h(v)^{p/2}.$$

Averaging over h and applying Lemma 9 yields the first line.

Applying Lemma 6 to the polynomial $Q_h(x) = \|vD(x)A(h)^T\|_4^4$, we find that for fixed h,

$$\mathop{\mathbb{E}}_x[\left\|vD(x)A(h)^T\right\|_4^{4p}] \leq O(p)^{2p} h(v)^p.$$

Taking an expectation over h and applying Lemma 9, we get that

$$\mathbb{E}[\left\|vD(x)A(h)^T\right\|_4^{4p}] \leq O(p)^{4p}\left(\frac{\|v\|_2^4}{m}\right)^{p} + O(p)^{4p}\|v\|_4^{4p}.$$

This completes the proof. □

We are now prepared to prove Lemma 28.

Proof of Lemma 28. We proceed by induction on i proving that the desired inequalities hold with probability at least $1 - i(\delta/n)^2$. As a base case we consider i so that $\|v\|_0^3 \leq n_i$. In this case, by repeated application of Lemma 25, we find that with at least the desired probability that $\|v\|_0 - \|v_i\|_0 \leq i\log(n/\delta)$. This implies that other than its zero coefficients, v_i has $\|v\|_0 - 2i\log(n/\delta)$ coefficients of norm 1, and at most $i\log(n/\delta)$ other coefficients each of norm at most $i\log(n/\delta)$. This means that

$$\|v_i\|_2^2 = \|v\|_2^2 + O(i^3\log^3(n/\delta)), \quad \text{and} \quad \|v_i\|_4^4 = \|v\|_2^2 + O(i^5\log^5(n/\delta)).$$

Our bounds follow immediately.

Otherwise, for $\|v_0\|^6 > n_i$, we proceed by induction on i. As a base case, note that the desired inequalities hold for $i = 1$ as $v_1 = v$, and $\|v\|_4 = \sqrt{\|v\|_2}$. We claim that if $\|v_i\|$ satisfies the desired inequalities, then v_{i+1} also does with probability at least $1 - (\delta/n)^2$. Note that $v_{i+1} = v_i D(Z_i)A(h_i)^T$. Note also that Z_i and h_i are k-wise independent for $k = C\log(n/\delta)/\log(n_i)$. Applying Lemma 29 with $p = \lfloor k/4 \rfloor$, we find that

$$\mathbb{E}\left[\left(\|v_i\|_2^2 - \|v_{i+1}\|_2^2\right)^{2p}\right] \leq O(p)^{4p}\left(\frac{\|v_i\|_2^4}{n_{i+1}}\right)^p + O(p)^{4p}\|v_i\|_4^{4p},$$

and

$$\mathbb{E}[\|v_{i+1}\|_4^{4p}] \leq O(p)^{4p}\left(\frac{\|v_i\|_2^4}{n_{i+1}}\right)^p + O(p)^{4p}\|v_i\|_4^{4p}.$$

Applying the Markov bound to the first of these equations we find that the probability that $\|v_{i+1}\|_2^2 \geq \|v_i\|_2^2 + 4^i\|v\|_2^2$ is at most

$$\begin{aligned}\left(\frac{O(p^4)}{n_{i+1}}\right)^p + O\left(\frac{p\|v_i\|_4}{\|v\|_2}\right)^{4p} &\leq n_{i+1}^{-p/2} + O(pn_{i+1}^{-1/10})^{4p} \\ &\leq n_{i+1}^{-p/2} + n_{i+1}^{-p/11} \\ &\leq (\delta/n)^2/2.\end{aligned}$$

Where the first inequality above is by the inductive hypothesis. This implies that $\|v_{i+1}\|_2 \leq 2^{i+1}\|v\|_2$ with the desired probability.

Applying the Markov bound to the latter of these equations we find that the probability that $\|v_{i+1}\|_4 > \|v\|_2/n_{i+1}^{1/20}$ is at most

$$O\left(\frac{p^4\|v_i\|_2^4}{\|v\|_2^4 n_{i+1}^{4/5}} + \frac{p^4\|v_i\|_4^4 n_{i+1}^{1/5}}{\|v_2\|_2^4}\right)^p \leq O\left(n_i^{-1/2} + \frac{n_{i+1}^{1/5}}{n_i^{1/5}}\right)^p \leq O(n_i^{-1/10})^p \leq (\delta/n)^2/2.$$

Where above we use that

$$\|v_i\|_4 \leq \frac{\|v\|_2}{\min(\|v\|_2^{1/3}, n_i^{1/20})} = \frac{\|v\|_2}{n_i^{1/20}}.$$

Thus, with the desired probability $\|v_{i+1}\|_4 \leq \|v\|_2/n_i^{1/20}$. This completes the inductive step, and finishes the proof. □

4.2.4 Combined analysis

We are now ready to prove Lemma 23.

Proof of Lemma 23. First note that if $i = t - 1$, the lemma follows immediately from the pseudorandomness properties of Z. We thus consider only $i < t - 1$.

We note that $v \cdot Y_i = v_i \cdot X_i$ and $v \cdot Y_{i+1} = v_i \cdot X_{i+1} A(h_i) D(Z_i)$. If $\|v\|_0^3 \leq n_{i+1}$, we are done by Lemma 24. Otherwise, assume that $\|v\|_0^3 > n_{i+1}$. By Lemma 28 we have that except for an event of probability $\delta/(4n)$ we have that

$$\|v_i\|_2 \leq \log(n)\|v\|_2 \quad \text{and} \quad \|v_i\|_4 \leq \|v\|_2 n_i^{-1/20}.$$

By ignoring the possibility that these are violated, we introduce an error of at most $\delta/(2n)$, thus it suffices to only consider the case where the choice of $h_1, Z_1, \ldots, h_{i-1}, Z_{i-1}$ are such that the above holds. We now need to bound

$$\left|\mathbb{E}[\phi_{v_i,\alpha}(X_i)] - \mathbb{E}[\phi_{v_i,\alpha}(X_{i+1}A(h_i)D(Z_i))]\right|.$$

Since X_i has the same distribution as $X_{i+1}A(h_i)D(X_i)$, we may apply Lemma 26 that for $p = \Omega\left(\frac{C\log(n/\delta)}{\log(n_i)}\right)$ that the above is bounded by

$$\begin{aligned} O(p)^{2p}\left(\frac{\alpha^4\|v_i\|_2^4}{n_{i+1}} + \alpha^4\|v_i\|_4^4\right)^{p/8} &\leq O(p)^{2p}\left(2^{4i}\log^{12}(1/\delta)n_i^{-1/2} + \log^{12}(1/\delta)n_i^{-1/5}\right)^{p/8} \\ &\leq \left(\log^{52}(n/\delta)n_i^{-1/5}\right)^{p/8} \\ &\leq n_i^{-p/50} \\ &\leq \delta/(2n). \end{aligned}$$

This completes the proof. □

Proposition 22 now follows immediately after noting that

$$\left|\mathbb{E}[\phi_{v,\alpha}(X)] - \mathbb{E}[\phi_{v,\alpha}(Y)]\right| \leq \sum_{i=1}^{t-1}\left|\mathbb{E}[\phi_{v,\alpha}(Y_i)] - \mathbb{E}[\phi_{v,\alpha}(Y_{i+1})]\right|.$$

4.2.5 Approximating a product of cosines

Here we prove Lemma 27.

Proof of Lemma 27. Note that so long as $\alpha S_i < 1/10$ for all i that by Taylor expansion we have that

$$\begin{aligned}
&\prod_{i=1}^{m} \cos(2\pi\alpha S_i) \\
&= \exp\left(-2\pi^2\alpha^2 \sum_{i=1}^{m} S_i^2 + \sum_{j=2}^{p/2-1} \left(c_j \sum_{i=1}^{m} (\alpha S_i)^{2j}\right) + \sum_{i=1}^{m} O(\alpha S_i)^p\right) \\
&= \exp(-2\pi^2\alpha^2 T) \exp\left(-2\pi^2\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right) + \sum_{j=2}^{p/2-1} \left(c_j \sum_{i=1}^{m} (\alpha S_i)^{2j}\right) + \sum_{i=1}^{m} O(\alpha S_i)^p\right)
\end{aligned}$$

where the c_j are constants obtained from the Taylor expansion of $\log(\cos(z))$. Furthermore, since $\log(\cos(z))$ is analytic in a disk around $z = 0$, we have that $c_j = O(1)^j$. Note by conditioning on whether or not $\sum_{i=1}^{m} O(\alpha S_i)^p$ is more than 1, we find that the above is equal to

$$\exp(-2\pi^2\alpha^2 T) \exp\left(-2\pi^2\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right) + \sum_{j=2}^{p/2-1} \left(c_j \sum_{i=1}^{m} (\alpha S_i)^{2j}\right)\right) \left(1 + \sum_{i=1}^{m} O(\alpha S_i)^p\right) + \sum_{i=1}^{m} O(\alpha S_i)^p.$$

For each j, let p_j be the ceiling of $p/(2j)$. Note that $p \le 2j \cdot p_j \le 2p$. Under the additional assumption that $\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right), \sum_{i=1}^{m} (\alpha S_i)^4 < a$, for some sufficiently small constant a we have that the above is equal to

$$\begin{aligned}
&\exp(-2\pi^2\alpha^2 T) \cdot \left(\sum_{t=0}^{p_2-1} \frac{\left(-2\pi^2\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right)\right)^t}{t!}\right) \cdot \left(1 + O\left(\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right)\right)^{p_2}\right) \\
&\quad \cdot \prod_{j=2}^{p/2-1} \left(\sum_{t=0}^{p_j-1} \frac{\left(c_j \sum_{i=1}^{m} (\alpha S_i)^{2j}\right)^t}{t!}\right) \cdot \left(1 + O(1)^p \left(\sum_{i=1}^{m} (\alpha S_i)^{2j}\right)^{p_j}\right) \\
&\quad \cdot \left(1 + \sum_{i=1}^{m} O(\alpha S_i)^p\right) + \sum_{i=1}^{m} O(\alpha S_i)^p \\
= \exp(-2\pi^2\alpha^2 T) &\cdot \left(\sum_{t=0}^{p_2-1} \frac{\left(-2\pi^2\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right)\right)^t}{t!}\right) \cdot \prod_{j=2}^{p/2-1} \left(\sum_{t=0}^{p_j-1} \frac{\left(c_j \sum_{i=1}^{m} (\alpha S_i)^{2j}\right)^t}{t!}\right) \\
&\quad + O(1)^p \left(\left(\alpha^2 \left(\sum_{i=1}^{m} S_i^2 - T\right)\right)^{p_2} + \left(\sum_{i=1}^{m} (\alpha S_i)^4\right)^{p/4} + \sum_{i=1}^{m} (\alpha S_i)^p\right)
\end{aligned}$$

Next consider the above term

$$\prod_{j=2}^{p/2-1}\left(\sum_{t=0}^{p_j-1}\frac{\left(c_j\sum_{i=1}^m(\alpha S_i)^{2j}\right)^t}{t!}\right).$$

Let it equal $P(S_i)+E(S_i)$ where P is the polynomial consisting of all the terms of total degree at most p. We note that for any j that $c_j\sum_{i=1}^m(\alpha S_i)^{2j}$ is at most

$$O(a)^{j/4}\left(\sum_{i=1}^m(\alpha S_i)^4\right)^{j/4}.$$

Therefore, $|E(S_i)|$ is at most $\left(\sum_{i=1}^m(\alpha S_i)^4\right)^{p/8}$ times the sum of the degree more than p coefficients in the Taylor expansion of

$$\exp\left(\frac{1}{1-O(a^{1/8}z)}\right).$$

For a sufficiently small, the above has radius of convergence more than 1, and thus the sum of the degree more than p terms is bounded. Thus, $E(S_i)$ is

$$O\left(\sum_{i=1}^m(\alpha S_i)^4\right)^{p/8}.$$

Therefore, assuming that $\alpha S_i<1/10$ for all i, and $\alpha^2\left(\sum_{i=1}^m S_i^2-T\right),\sum_{i=1}^m(\alpha S_i)^4<a$, then

$$\prod_{i=1}^m\cos(2\pi\alpha S_i)$$

equals

$$\exp(-2\pi^2\alpha^2T)\cdot\left(\sum_{t=0}^{p_2-1}\frac{\left(-2\pi^2\alpha^2\left(\sum_{i=1}^m S_i^2-T\right)\right)^t}{t!}\right)P(S_i)$$
$$+O(1)^p\left(\left(\alpha^2\left(\sum_{i=1}^m S_i^2-T\right)\right)^{p_2}+\left(\sum_{i=1}^m(\alpha S_i)^4\right)^{p/8-1}+\sum_{i=1}^m(\alpha S_i)^p\right).$$

On the other hand, if the stated assumptions fail, the main term above is bounded by a polynomial in $\alpha^2\left(\sum_{i=1}^m S_i^2-T\right)$ and $\sum_{i=1}^m(\alpha S_i)^4$ with total degree at most $2p$ and sum of coefficients $O(1)^p$. Therefore, under no additional assumptions we have that

$$\prod_{i=1}^m\cos(2\pi\alpha S_i)$$

equals

$$\exp(-2\pi^2\alpha^2 T)\cdot\left(\sum_{t=0}^{p_2-1}\frac{\left(-2\pi^2\alpha^2\left(\sum_{i=1}^m S_i^2-T\right)\right)^t}{t!}\right)P(S_i)$$
$$+O(1)^p\left(\left(\alpha^2\left(\sum_{i=1}^m S_i^2-T\right)\right)^{p/2}+\left(\alpha^2\left(\sum_{i=1}^m S_i^2-T\right)\right)^{p}+\left(\sum_{i=1}^m(\alpha S_i)^4\right)^{p/8}+\left(\sum_{i=1}^m(\alpha S_i)^4\right)^{p/2}\right).$$

The claim now follows. □

4.3 Final analysis

We can finally state our main generator and prove Theorem 2.

Proof of Theorem 2. Let Y_1, Y_2 be the generators from Propositions 19 and 22 for $\delta = \epsilon/6n$. Let Y be the the coordinate-wise product of the strings Y_1, Y_2. We claim that for any $v \in \{-1,0,1\}^n$ and $X \in_u \{\pm 1\}^n$,

$$d_{TV}(v\cdot X, v\cdot Y) \le \epsilon. \tag{19}$$

The theorem follows immediately from the above claim and the bounds on the seed-lengths from Propositions 19 and 22.

To prove the theorem, we first prove that for all $\alpha \in \mathbb{R}$,

$$|\mathbb{E}[\phi_{v,\alpha}(X)] - \mathbb{E}[\phi_{v,\alpha}(Y)]| \le \epsilon/(2n).$$

Now, if $\log^3(1/\delta)/\left\|v\right\|_2 \le \alpha$, then

$$\mathbb{E}[\phi_{v,\alpha}(Y)] = \mathbb{E}[\phi_{D(Y_2)v,\alpha}(Y_1)]$$

and

$$\mathbb{E}[\phi_{v,\alpha}(X)] = \mathbb{E}[\phi_{v,\alpha}(D(Y_2)X)] = \mathbb{E}[\phi_{D(Y_2)v,\alpha}(X)].$$

However by Proposition 19, we have that

$$\left|\mathbb{E}[\phi_{D(Y_2)v,\alpha}(Y_1)] - \mathbb{E}[\phi_{D(Y_2)v,\alpha}(X)]\right| \le \epsilon/(3n).$$

Similarly, if $\alpha \le \log^3(1/\delta)/\left\|v\right\|_2$, then then note that

$$\mathbb{E}[\phi_{v,\alpha}(Y)] = \mathbb{E}[\phi_{D(Y_1)v,\alpha}(Y_2)]$$

and

$$\mathbb{E}[\phi_{v,\alpha}(X)] = \mathbb{E}[\phi_{v,\alpha}(D(Y_1)X)] = \mathbb{E}[\phi_{D(Y_1)v,\alpha}(X)].$$

However by Proposition 22, we have that

$$\left|\mathbb{E}[\phi_{D(Y_1)v,\alpha}(Y_2)] - \mathbb{E}[\phi_{D(Y_1)v,\alpha}(X)]\right| \leq \epsilon/(3n).$$

Thus, we have our result for all $\alpha \in [0, 1/4]$. Noting that $\phi_{v,-\alpha}(X) = \overline{\phi_{v,\alpha}(X)}$, we determine that the statement in question holds for α if and only if it holds for $-\alpha$. Thus, the inequality in question holds for all $\alpha \in [-1/4, 1/4]$. Next, note that for any $X \in \{\pm 1\}^n$ that $\phi_{v,\alpha+1/2}(X) = \exp(\pi i v \cdot X)\phi_{v,\alpha}(X) = (-1)^{\|v\|_0}\phi_{v,\alpha}(X)$. Thus, the statement in question holds for α if and only if it holds for $\alpha + 1/2$. Thus, it holds for all real α. Equation 19 now follows from the above argument and Claim 2 applied to $Z_1 = \langle v, X\rangle$ and $Z_2 = \langle v, Y\rangle$. □

References

[BO10] Eric Blais and Ryan O'Donnell. Lower bounds for testing function isomorphism. In *Computational Complexity (CCC), 2010 IEEE 25th Annual Conference on*, pages 235–246. IEEE, 2010.

[BRRY14] Mark Braverman, Anup Rao, Ran Raz, and Amir Yehudayoff. Pseudorandom generators for regular branching programs. *SIAM J. Comput.*, 43(3):973–986, 2014.

[BV10] Joshua Brody and Elad Verbin. The coin problem and pseudorandomness for branching programs. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 30–39, 2010.

[CRSW13] L. Elisa Celis, Omer Reingold, Gil Segev, and Udi Wieder. Balls and bins: Smaller hash families and faster evaluation. *SIAM J. Comput.*, 42(3):1030–1050, 2013.

[De11] Anindya De. Pseudorandomness for permutation and regular branching programs. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity, CCC 2011, San Jose, California, June 8-10, 2011*, pages 221–231, 2011.

[De14] Anindya De. Beyond the central limit theorem: asymptotic expansions and pseudorandomness for combinatorial sums, 2014. ECCC, TR14-125.

[DGJ+09] Ilias Diakonikolas, Parikshit Gopalan, Ragesh Jaiswal, Rocco A. Servedio, and Emanuele Viola. Bounded independence fools halfspaces. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS '09)*, 2009.

[DKN10] Ilias Diakonikolas, Daniel Kane, and Jelani Nelson. Bounded independence fools degree-2 threshold functions. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS '10)*, 2010.

[Fel71] William Feller. *An Introduction to Probability Theory and Its Applications, Vol. 2 (Volume 2)*. Wiley, New York, London, Sydney, 2nd edition, 1971.

[GMR+12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Better pseudorandom generators from milder pseudorandom restrictions. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 120–129, 2012.

[GMRZ13] Parikshit Gopalan, Raghu Meka, Omer Reingold, and David Zuckerman. Pseudorandom generators for combinatorial shapes. *SIAM J. Comput.*, 42(3):1051–1076, 2013.

[GNR14] Parikshit Gopalan, Noam Nisan, and Tim Roughgarden, 2014. Manuscript in preparation.

[GOWZ10] Parikshit Gopalan, Ryan O'Donnell, Yi Wu, and David Zuckerman. Fooling functions of halfspaces under product distributions. In *25th Annual IEEE Conference on Computational Complexity*, pages 223–234, 2010.

[GR09] Parikshit Gopalan and Jaikumar Radhakrishnan. Finding duplicates in a data stream. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 402–411, 2009.

[HKM12] Prahladh Harsha, Adam Klivans, and Raghu Meka. An invariance principle for polytopes. *J. ACM*, 59(6):29, 2012.

[INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 356–364, 1994.

[IW97] Russell Impagliazzo and Avi Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *STOC*, pages 220–229, 1997.

[Kan11a] Daniel M. Kane. k-independent gaussians fool polynomial threshold functions. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity, CCC 2011, San Jose, California, June 8-10, 2011*, pages 252–261, 2011.

[Kan11b] Daniel M. Kane. A small PRG for polynomial threshold functions of gaussians. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 257–266, 2011.

[Kan14] Daniel M. Kane. A pseudorandom generator for polynomial threshold functions of gaussian with subpolynomial seed length. In *IEEE 29th Conference on Computational Complexity, CCC 2014, Vancouver, BC, Canada, June 11-13, 2014*, pages 217–228, 2014.

[KM14] Pravesh Kothari and Raghu Meka. Almost-optimal pseudorandom generators for spherical caps, 2014. Manuscript.

[KMN11] Daniel M. Kane, Raghu Meka, and Jelani Nelson. Almost optimal explicit johnson-lindenstrauss families. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 14th International Workshop, APPROX 2011, and 15th International Workshop, RANDOM 2011, Princeton, NJ, USA, August 17-19, 2011. Proceedings*, pages 628–639, 2011.

[KNP11] Michal Koucký, Prajakta Nimbhorkar, and Pavel Pudlák. Pseudorandom generators for group products: extended abstract. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 263–272, 2011.

[KRS12] Zohar Shay Karnin, Yuval Rabani, and Amir Shpilka. Explicit dimension reduction and its applications. *SIAM J. Comput.*, 41(1):219–249, 2012.

[LRTV09] Shachar Lovett, Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Pseudorandom bit generators that fool modular sums. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, pages 615–630, 2009.

[MOO05] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS '05)*, pages 21–30, 2005.

[MORS09] Kevin Matulef, Ryan O'Donnell, Ronitt Rubinfeld, and Rocco A. Servedio. Testing ± 1-weight halfspace. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, pages 646–657, 2009.

[MZ09] Raghu Meka and David Zuckerman. Small-bias spaces for group products. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, pages 658–672, 2009.

[MZ13] Raghu Meka and David Zuckerman. Pseudorandom generators for polynomial threshold functions. *SIAM J. Comput.*, 42(3):1275–1301, 2013.

[Nis92] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.

[Nis94] Noam Nisan. $\mathsf{RL} \subseteq \mathsf{SC}$. *Computational Complexity*, 4(1):1–11, 1994.

[NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. on Comput.*, 22(4):838–856, 1993.

[NZ96] Noam Nisan and David Zuckerman. Randomness is linear in space. *J. Comput. System Sci.*, 52(1):43–52, 1996.

[O'D14] Ryan O'Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014.

[Rei08] Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4), 2008.

[RR99] Ran Raz and Omer Reingold. On recycling the randomness of states in space bounded computation. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999, Atlanta, Georgia, USA*, pages 159–168, 1999.

[RS10] Yuval Rabani and Amir Shpilka. Explicit construction of a small epsilon-net for linear threshold functions. *SIAM J. Comput.*, 39(8):3501–3520, 2010.

[RS13] D. Ron and R. Servedio. Exponentially improved algorithms and lower bounds for testing signed majorities. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1319–1336, 2013.

[RTV06] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Pseudorandom walks on regular digraphs and the RL vs. L problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 457–466, 2006.

[SSS95] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2):223–250, 1995.

[SZ99] Michael E. Saks and Shiyu Zhou. BP $_{\mathrm{h}}$space(s) subseteq dspace(s$^{3/2}$). *J. Comput. Syst. Sci.*, 58(2):376–403, 1999.

[VW08] Emanuele Viola and Avi Wigderson. Norms, xor lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4(7):137–168, 2008.

A Proofs from Section 2

Proof of Lemma 7. This follows from the fact that $\|Q^p\|_1 \leq \|Q\|_1^p$. Therefore,

$$\mathbb{E}\left[Q(x)^p\right] \leq \mathop{\mathbb{E}}_{X \in_u \{\pm 1\}^n}[Q(X)^p] + \|Q\|_1^p \delta \leq (p-1)^{pd/2} \cdot \|Q\|_2^p + \|Q\|_1^p \delta.$$

□

Proof of Lemma 8. Note that because $\|v\|_2 = 1$ that $\|v\|_1 \leq \|v\|_0^{1/2}$ by Cauchy-Schwarz. We note by the Markov inequality that for even p that

$$\Pr[|\langle v,\ x\rangle| > t] \leq t^{-p}\, \mathbb{E}[|\langle v,\ x\rangle|^p].$$

We need a slightly strengthened version of Lemma 7 to bound this. Note that if $f(x) = \langle v,\ x\rangle$

$$\mathbb{E}[f(x)] \leq \|f^p\|_0 \epsilon + \|f\|_p^p \leq \|v\|_0^p \epsilon + (p-1)(p-3)\cdots 1.$$

The bound on $\|f\|_p$ comes from noting that the expectation of f^p under Gaussian inputs is $(p-1)(p-3)\cdots 1$ and that the expectation under Bernoulli inputs is at most this (which can be seen by expanding and comparing terms). Therefore, we have that

$$\Pr[|\langle v,\ x\rangle| > t] \leq t^{-p}\sqrt{2}(p/e)^{p/2} + t^{-p}\|v\|_0^p \epsilon$$

Letting p be the largest even integer less than t^2, we find that this is at most

$$\sqrt{2}\exp(-p/2) + \|v\|_0^{t^2}\epsilon,$$

which is sufficient when $t \geq 2$. For $1 \leq t \leq \sqrt{2}$, the trivial upper bound of 1 is sufficient, and for $\sqrt{2} \leq t \leq 2$, we may instead use the bound for $p = 2$. □

Proof of Lemma 9. Let $I_{i,k}$ be the indicator function of the event that $h(i) = k$. Note that $h(v) = \sum_{i,j,k} I_{i,k} I_{j,k} v_i^2 v_j^2$. Therefore,

$$h(v)^p = \sum_{i_1,\ldots,i_p,j_1,\ldots,j_p} \sum_{k_1,\ldots,k_p} \prod_{t=1}^{p} I_{i_t,k_t} I_{j_t,k_t} \prod_{t=1}^{p} v_{i_t}^2 v_{j_t}^2.$$

Let $R(i_t, j_t, k_t)$ be 0 if for some t, t' $k_t \neq k'_t$ but one of i_t or j_t equals $i_{t'}$ or $j_{t'}$ and otherwise be equal to m^{-T} where T is the number of distinct values taken by i_t or j_t. Notice that by the δ-biasedness of h that

$$\mathbb{E}\left[\prod_{t=1}^{p} I_{i_t,k_t} I_{j_t,k_t}\right] \leq R(i_t, j_t, k_t) + \delta.$$

Combining with the above we find that

$$\begin{aligned}
\mathbb{E}[h(v)^p] &\leq \sum_{i_1,\ldots,i_p,j_1,\ldots,j_p} \sum_{k_1,\ldots,k_p} (R(i_t, j_t, k_t) + \delta) \prod_{t=1}^{p} v_{i_t}^2 v_{j_t}^2 \\
&\leq \sum_{i_1,\ldots,i_p,j_1,\ldots,j_p} \sum_{k_1,\ldots,k_p} R(i_t, j_t, k_t) \prod_{t=1}^{p} v_{i_t}^2 v_{j_t}^2 + \delta m^p \sum_{i_1,\ldots,i_p,j_1,\ldots,j_p} \prod_{t=1}^{p} v_{i_t}^2 v_{j_t}^2 \\
&\leq \sum_{i_1,\ldots,i_p,j_1,\ldots,j_p} \sum_{k_1,\ldots,k_p} R(i_t, j_t, k_t) \prod_{t=1}^{p} v_{i_t}^2 v_{j_t}^2 + \delta m^p \|v\|_2^{4p}.
\end{aligned}$$

Next we consider

$$\sum_{k_1,\ldots,k_p} R(i_t, j_t, k_t)$$

for fixed values of $i_1, \ldots, i_p, j_1, \ldots, j_p$. We claim that it is at most $m^{-S/2}$ where S is again the number of distinct elements of the form i_t or j_t that appear in this way an odd number of times. Letting T be the number of distinct elements of the form i_t or j_t, the expression in question is m^{-T} times the number of choices of k_t so that each value of i_t or j_t appears with only one value of k_t. In other words this is m^{-T} times the number of functions $f : \{i_t, j_t\} \to [m]$ so that $f(i_t) = f(j_t)$ for all t. This last relation splits $\{i_t, j_t\}$ into equivalence classes given by the transitive closure of the operation that $x \sim y$ if $x = i_t$ and $y = j_t$ for some t. We note that any x that appears an odd number of times as an i_t or j_t must be in an equivalence class of size at least 2 because it must appear at least once with some other element. Therefore, the number of equivalence classes, E is at least $T - S/2$. Thus, the sum in question is at most $m^{-T} m^E \leq m^{-S/2}$. Therefore, we have that

$$\mathbb{E}[h(v)^p] \leq (2p)! \sum_{\text{Multisets } M \subset [n], |M| = 2p} m^{-\{\text{Odd}(M)\}/2} \prod_{i \in M} v_i^2 + \delta m^p \|v\|_2^{4p}.$$

Where $\text{Odd}(M)$ is the number of elements occurring in M an odd number of times. This

equals

$$\begin{aligned}
\mathbb{E}[h(v)^p] &\leq (2p)! \sum_{k=0}^{p} \sum_{\text{Multisets } M \subset [n], |M|=2p, \text{Odd}(M)=2k} m^{-k} \prod_{i \in M} v_i^2 + \delta m^p \|v\|_2^{4p} \\
&\leq (2p)! \sum_{k=0}^{p} m^{-k} \sum_{i_1,\ldots,i_{2k}} \sum_{j_1,\ldots,j_{p-k}} \prod v_{i_t}^2 \prod v_{j_t}^4 + \delta m^p \|v\|_2^{4p} \\
&= (2p)! \sum_{k=0}^{p} \left(\frac{\|v\|_2^4}{m} \right)^k \|v\|_4^{4(p-k)} + \delta m^p \|v\|_2^{4p} \\
&\leq O(p)^{2p} \left(\frac{\|v\|_2^4}{m} \right)^p + O(p)^{2p} \|v\|_4^{4p} + \delta m^p \|v\|_2^{4p}.
\end{aligned}$$

Note that the second line above comes from taking M to be the multiset

$$\{i_1, i_2, \ldots, i_{2k}, j_1, j_1, j_2, j_2, \ldots, j_{p-k}, j_{p-k}\}.$$

This completes our proof. □

Proof of Lemma 10. Let X_i denote the indicator random variable which is 1 if $h(i) = j$ and 0 otherwise. Let $Z = \sum_i v_i X_i$. Now, if h were a truly random hash function, then, by Hoeffding's inequality,

$$\Pr\left[\left|Z - \|v\|_1 / m\right| \geq t\right] \leq 2 \exp\left(-t^2 / 2 \sum_i v_i^2 \right).$$

Therefore, for a truly random hash function and even integer $p \geq 2$, $\|Z\|_p = O(\|v\|_2)\sqrt{p}$. Therefore, for a δ-biased hash family, we get $\|Z\|_p^p \leq O(p)^{p/2} \|v\|_2^p + \|v\|_1^p \delta$. Hence, by Markov's inequality, for any $t > 0$,

$$\Pr\left[\left|Z - \|v\|_1 / m\right| \geq t\right] \leq \frac{O(p)^{p/2} \|v\|_2^p + \|v\|_1^p \delta}{t^p}.$$

□